How to Keep a Lighthouse Honest: AIPAC and Raising the Cost of Workarounds in Modern Political Influence

Douglas C. Youvan

doug@youvan.com

www.youvan.ai

October 9, 2025

When explicit quid-pro-quo is rare and everyone "knows the checklist," the smartest move isn't to hunt for confessions—it's to re-engineer the harbor. This paper reframes contemporary political influence (PACs, super-PACs, donor networks, coordinated "independent" operations) as an adversarial systems problem. We distinguish lawful pressure from illegal bribery/extortion, then focus on what actually bites when actors optimize around the rules: instrumentation, chokepoints, and narrative accountability. Drawing on an engineering mindset, we propose an evidence stack—creative-hash ad similarity, vendor-lattice analysis, impression-timing around floor action, language-overlap diffing—paired with institutional redesign (top-four + RCV primaries, democracy vouchers, rapid disclosure APIs, exclusive-vendor rules, cooling-off windows, creative-provenance tags). We add litigation and enforcement wedges (firewall fraud, conduit schemes, platform archives), whistleblower channels that protect junior staff, and public dashboards that translate forensics into plain signals. A companion poetic metaphor—the lighthouse—keeps the program human: when fog is policy and currents are money, the goal isn't purity; it's making the cheapest path the honest one. The thesis is simple: if you can't prove the bargain, price the workaround. Either the trick becomes impossible, or it grows teeth and bites its owner. This is a design brief for citizens, journalists, reformers, and principled candidates.

Keywords: campaign finance, independent expenditures, coordination detection, political coercion, adversarial auditing, vendor lattice, ad hash similarity, language overlap, ranked choice voting, top four primaries, democracy vouchers, transparency APIs, whistleblower protection, conduit donations, firewall fraud, platform archives, litigation strategy, narrative accountability, lighthouse protocol.

A collaboration with GPT-5. 68 pages. CC4.0.

Note: The Epilogue compares this approach with John Nash's economics model.

Outline

1. Framing the Problem
   1.1 From “catch the email” to “raise the cost”
   1.2 Legal vs. normative: pressure, bribery, extortion, compellence
2. A Minimal Vocabulary of Influence
   2.1 Donations, IEs, bundling, conduits
   2.2 Coordination and its paper walls
3. Threat Models and Actor Incentives
   3.1 How sophisticated campaigns evade the checklist
   3.2 The optimization loop: law → workaround → norm drift
4. Evidence Stack I: Adversarial Forensics
   4.1 Creative-hash similarity for TV/digital assets
   4.2 Impression-timing vs. floor action and markups
   4.3 Vendor-lattice mapping and sub-vendor leakage
   4.4 Message diffing: n-gram and script overlap windows
5. Evidence Stack II: Causal Inference Without Confessions
   5.1 Event studies around spending shocks
   5.2 Regression discontinuity in near-tie races
   5.3 Synthetic controls for policy flips
6. Chokepoint Design
   6.1 Exclusive-vendor/affiliate rules (covering subs)
   6.2 Cooling-off windows for staff transitions
   6.3 Creative-provenance tags and broadcaster/platform duties
7. Move the Battlefield
   7.1 Top-four + ranked-choice to dilute single-issue gatekeeping
   7.2 Democracy vouchers / small-donor matching as counter-leverage
   7.3 Fast-disclosure infrastructure: 24-hour IE & ad APIs
8. Litigation & Enforcement Wedges
   8.1 Firewall fraud and deceptive compliance
   8.2 Conduit schemes (LLCs, straw donors): bank/wire metadata
   8.3 State-level transparency and archive mandates
9. Whistleblowers Who Actually Matter
   9.1 Protecting junior staff and vendors
   9.2 Bounties, safe channels, audit-ready pledges

10. Public Dashboards & Narrative Tools
    10.1 Pay-to-Vote Index (outside-money shocks vs. roll-call flips)
    10.2 Race dashboards: ad-hash overlap, vendor-overlap meters, 48-hour IE tickers
    10.3 Translating forensics into citizen-readable signals

11. Ethics of Design: Making the Honest Path Cheap
    11.1 Minimizing false positives and chilling effects
    11.2 Due process for accused campaigns and PACs

12. Case Typology (Abstracted)
    12.1 “Clone creatives” pattern
    12.2 “Common agent” pattern
    12.3 “Timed barrage” pattern
    12.4 “Conduit carousel” pattern

13. The Lighthouse Protocol (Poetic Interlude)
    13.1 Metaphor as design spec: beams, fog, currents, wakes
    13.2 Instrument the shoreline; don’t sermonize the sea

14. Implementation Roadmap
    14.1 Data standards, IDs, upload flows, and public APIs
    14.2 Pilot jurisdictions and evaluation metrics
    14.3 Maintenance: audits, red-team drills, iteration

15. Limits and Failure Modes
    15.1 Adaptive adversaries and creative lawyering
    15.2 Jurisdictional patchwork and platform resistance
    15.3 Measuring impact without counterfactual omniscience

16. Conclusion — Pricing the Workaround
    16.1 Either the trick becomes impossible, or it grows teeth
    16.2 A citizen’s, journalist’s, and candidate’s to-do list

Appendix A: “Coordination Cost Raiser” One-Pager (field checklist)
Appendix B: Data Schemas for Ad Hashes, Vendor Graphs, and Disclosure APIs
Appendix C: Sample Statutes/Rules (exclusive vendors, provenance tags, archives)
Appendix D: Methods Notes (event study, RD, synthetic control templates)
Appendix E: “How to Keep a Lighthouse Honest” — full poem (for public comms)

Epilogue

References

1. Framing the Problem

1.1 From “catch the email” to “raise the cost”

Traditional anti-corruption strategy hunts for explicit quid-pro-quo: the email, the call, the text that says “vote X for $Y.” Sophisticated actors have adapted. They avoid explicit asks, route signals through intermediaries, and exploit lawful channels (independent expenditures, scorecards, primary threats) that discipline officials without ever stating the trade. Chasing confessions becomes a low-yield game.

The systems approach flips the objective:

- **Instrument the environment, not the villain.** If coordination leaves statistical, linguistic, and operational fingerprints (creative reuse, synchronized timing, vendor overlap), build detectors for those fingerprints.
- **Target chokepoints.** Make it difficult to share creatives, staff, and data across “independent” entities; require provenance tags and exclusive-vendor rules; shorten disclosure windows so “dark” tactics lose first-mover advantage.
- **Shift incentives.** Election design (top-four + RCV) and public financing (vouchers/matching) dilute the marginal power of a single funder’s threat.
- **Translate complexity into public signals.** Dashboards (ad-hash overlap meters, 48-hour IE tickers, Pay-to-Vote indices) convert forensics into reputational and electoral costs.

Design target: **price the workaround**. If explicit quid-pro-quo is rare, the goal is to make stealth coordination riskier, more expensive, and less decisive—so the cheapest, safest path is the honest one.

Operational test

1. Can a tactic be executed without leaving cross-entity artifacts (assets, vendors, timing)?
2. If artifacts exist, are they machine-detectable quickly enough to matter to voters/regulators?
3. Do rules and market design ensure that exposure triggers real costs (legal, reputational, electoral)?

If the answer to (2) or (3) is “no,” raise the cost.

1.2 Legal vs. normative: pressure, bribery, extortion, compellence

Precision in language keeps reforms sharp and defensible.

Legal categories (U.S.-centric)

- **Bribery:** an **explicit quid-pro-quo**—money or value **for** an **official act**. High proof bar; protected "access" and vague promises are typically insufficient.
- **Extortion (Hobbs Act):** obtaining property via force, fear, or "under color of official right." Threat must be unlawful (or the taking is wrongful).
- **Illegal coordination:** "Independent" spend that is in fact planned with a campaign (shared strategy, creatives, vendors, materials).
- **Conduit/Straw donations:** disguising the true source of funds.

Normative/analytical categories (political science)

- **Pressure/discipline:** lawful but forceful tactics—fund (or fund against) based on votes, publish scorecards, mobilize primaries.
- **Compellence:** "Do X or we impose costs" (distinct from deterrence: "Don't do Y or else"). States use compellence via sanctions, embargoes, energy cutoffs; domestic analogs appear in campaign finance via credible threats of heavy independent expenditures.

Why the distinction matters

- **Legal claims require evidence of prohibited elements** (explicit exchange, unlawful threats, coordination). Without that, assertions of "bribery/extortion" will fail in court and can backfire politically.
- **Normative claims guide design.** Even when conduct is legal, societies may judge it corrosive and tighten rules (e.g., faster disclosure, vendor firewalls, provenance requirements) or redesign markets (RCV, vouchers) to blunt its power.

Decision tree for classification

1. **Is there an explicit exchange tied to an official act?**
   – Yes → likely **bribery**.
2. **Is there a threat of unlawful harm or wrongful taking?**
   – Yes → likely **extortion**.

3. **Are nominally independent actors sharing creatives/data/staff/strategy?**
   – Evidence says yes → **illegal coordination**.
4. **None of the above, but large money is used to punish/reward votes?**
   – **Normative pressure/compellence**: lawful influence; address via disclosure, chokepoints, and election design.

Policy stance

- Reserve criminal labels for provable crimes.
- Treat everything else as an **engineering problem**: minimize exploitability, surface signals quickly, and ensure that evasive maneuvers carry visible costs.

2. A Minimal Vocabulary of Influence

2.1 Donations, IEs, bundling, conduits

Direct contributions (to campaigns/parties).

- *What they are:* Money given to a candidate committee or party, capped and disclosed.
- *Mechanics:* Hard-dollar limits; donor identity disclosed; timing/reporting rules.
- *Leverage channel:* Access and attention (meetings, call time), not lawful quid-pro-quo.
- *Signals to watch:* Sudden spikes around key votes; clustered donors from a common industry; new PACs with identical treasurers.

Independent expenditures (IEs).

- *What they are:* Spending "independent" of a campaign to advocate for/against a candidate (ads, mailers, digital). Unlimited if truly independent.
- *Mechanics:* Super-PACs and 501(c)(4)/(6) vehicles often execute; disclosure varies by entity and medium.
- *Leverage channel:* Credible threat of large, rapid negative/positive spend in primaries or close generals.
- *Signals to watch:* Content mirroring campaign creatives; synchronized ad drops with floor action; repeated vendor overlaps across races.

Bundling.

- *What it is:* Intermediaries aggregate many small donations and deliver them as a package (sometimes with tracking to credit the bundler).
- *Mechanics:* Legal with disclosure; creates "relational capital" between networks and officeholders.
- *Leverage channel:* Gatekeeping to donor lists, recurring events, and timed pledges.
- *Signals to watch:* Bundler-led giving waves preceding committee assignments; identical memo lines/LLC sources; donor networks shifting en bloc after a policy signal.

Conduits and straw donors.

- *What they are:* Illicit disguises of the true source of money (reimbursed donors, shell LLCs, pass-through nonprofits).
- *Mechanics:* Criminal if used to evade limits or conceal the source; often paired with opaque entities.
- *Leverage channel:* Same as contributions/IEs—but with reduced traceability.
- *Signals to watch:* Fresh LLCs with no business history; donors giving beyond apparent means; circular wires among related entities.

Dark-bright continuum (opacity vs. transparency).

- *Bright money:* Named donors to hard-dollar accounts; high auditability.
- *Dim money:* Super-PAC IEs with donor disclosure but complex networks.
- *Dark money:* 501(c) IEs/issue ads with donor anonymity or laundered sources.
- *Design target:* Pull activity toward the bright end with faster filings, provenance tags, cross-media archives.

Non-monetary transfers (often overlooked).

- *Data & lists:* Voter files, look-alike audiences, micro-targeting models.
- *Talent & time:* Staff rotations, "informal" advice, shared contractors.
- *In-kind media:* Coordinated op-eds, influencer placements, content syndication.
- *Signals to watch:* Identical segmentations/creative briefs; "volunteer" consultants who invoice an allied nonprofit.

2.2 Coordination and its paper walls

Coordination (bright line in law, gray in practice).

- *Definition:* An IE ceases to be "independent" if it is made "in cooperation, consultation, or concert with, or at the request or suggestion of" a candidate or party.
- *Why it matters:* Crossing this line converts unlimited IE into an in-kind contribution—often illegal due to size/source limits.

The paper walls.

- *Firewall memos:* Firms serving both a campaign and an IE promise internal separation (teams, files, comms).
- *Public materials safe harbor:* IEs can legally use publicly available info (speeches, websites) without "coordination."
- *Non-contact doctrine:* No emails, no calls, no shared plans—on paper.

How real coordination happens anyway (playbook).

1. **Common-agent signaling:** Shared parent firm or affiliate "reads" a campaign's needs; no direct message required.
2. **Creative mirroring:** IE repurposes campaign scripts/layouts/theme within hours ("independently" derived from public signals).
3. **Calendar cues:** Campaign schedules a rally/press hit knowing the IE will time a spend to amplify.
4. **Intermediary whispers:** Party committees, consultants, or aligned nonprofits transmit "public" priorities that are practically private.
5. **Talent carousel:** Staff rotate across entities with cooling-off fig-leaves; institutional memory acts as the conduit.
6. **Data osmosis:** Look-alike models and segments migrate via vendors that serve "everyone," functionally syncing strategy.

Forensic fingerprints (to pierce paper walls).

- *Asset lineage:* Hash-level similarity of audio/video, matching B-roll, identical lower-thirds/VO cadence.
- *Temporal coupling:* IE bursts tightly bracketing committee markups, whip counts, or last-minute amendments.

- *Vendor lattice:* Sub-vendor overlaps (edit bays, colorists, DSP seats, cloud workspaces) beneath the "separate" prime.
- *Language overlap:* N-gram/syntax matches between campaign emails/pressers and IE ads within 24–72 hours.
- *Targeting twins:* Near-identical audience definitions and geos; same exclusion lists; same flighting rhythm.

Strengthening the walls (design levers).

- *Exclusive-vendor rules:* One corporate family can't touch both sides of a race (cover subs/affiliates/franchises).
- *Cooling-off periods:* Enforce real gaps for staff moving between campaigns and IEs affecting that race.
- *Provenance & archives:* Platform/broadcaster duty to publish creative asset IDs, uploaders, buys, and targeting slices in 24 hours.
- *Real-time disclosure:* API-first filings for IEs; machine-readable vendor IDs and contract metadata.
- *Public comparators:* Official ad libraries that auto-flag content/targeting clones and vendor overlaps.

Classification checklist (practical).

- *If* explicit planning/requests exist → **illegal coordination**.
- *If* no explicit planning, but creatives, timing, vendors, and targets align beyond chance → **presumptive coordination** (administrative/enforcement action; civil/criminal depending on jurisdiction).
- *If* only shared ideology and public materials explain similarity → **legal parallelism** (address via faster disclosure and incentive redesign, not prosecution).

Principle for the rest of the paper:
We do not chase whispered words; we **instrument the ecosystem**. The objective is not to prove a thought crime but to make covert coordination **operationally hard, quickly visible, and reputationally costly**.

3. Threat Models and Actor Incentives

3.1 How sophisticated campaigns evade the checklist

Threat surface. Assume capable actors with legal counsel, media shops, data science, and allied nonprofits/super-PACs. Objective: maximize electoral leverage while minimizing prosecutable artifacts.

Evasion patterns (with counter-signals):

- **No-ask choreography.** Campaigns never utter a request; instead they seed *public* cues (press hits, slogans, B-roll, color palettes, "issues pages").
  • *Counter-signal:* ad-hash lineage + sub-vendor overlaps reveal recycled creative DNA within 24–72 hours.
- **Compliance theater.** Thick firewall memos, separate LLCs, different letterheads—same parent, servers, or freelancers.
  • *Counter-signal:* cloud workspace forensics, font/package fingerprints, build metadata in exported assets.
- **Intermediary fog.** Party committees or 501(c)s "brief the world," knowing only one IE can act at scale on three hours' notice.
  • *Counter-signal:* temporal coupling of IE buys to committee markups/whip calls; repeated priority mirroring.
- **Talent carousel.** Staff "cool off" for token weeks, then join the IE; institutional memory substitutes for messages.
  • *Counter-signal:* HR graphing (titles, vendors, Slack/Git timestamps), identical audience taxonomies re-appearing post-move.
- **Data osmosis.** Shared vendors run look-alike models for multiple clients; seed features leak via "best practices."
  • *Counter-signal:* targeting-slice congruence (same geos, dayparts, exclusions), identical lift patterns across media.
- **LLC laundromats.** Pop-up entities obscure donors, consultants, or asset ownership; invoices routed through affiliates.
  • *Counter-signal:* bank/Secretary-of-State joins, registered agent reuse, wire-timing clusters.
- **Narrative synchronization.** Campaign staff speak in public "to the air"; IE ads answer within hours with matching frames.

• *Counter-signal:* n-gram/syntax overlap + audio VO cadence similarity + creative motif reuse (lower-thirds, bumpers).

Red-team model.

- *Attacker goal:* move votes at lowest legal risk.
- *Tools:* ambiguity, time compression, shared infrastructure, plausible public sourcing.
- *Weakest link:* production supply chain (assets, people, platforms) leaves fingerprints.
- *Defender goal:* make fingerprints fast, comparable, and costly.

3.2 The optimization loop: law → workaround → norm drift

Step 1: Law. Jurisdictions draw bright lines (no coordination; disclose donors; cap direct contributions). Courts narrow doctrines (e.g., tight "official act" definitions; broad protection for independent expenditures).

Step 2: Workaround. Actors route influence through what remains cheap and powerful: IEs, primaries, bundler networks, opaque nonprofits, creative reuse, common agents, staff mobility, and "public" signaling.

Step 3: Norm drift. Repetition reframes yesterday's workaround as today's standard practice. Vendors productize evasion (turnkey firewalls, shell LLC kits, creative libraries). Journalistic fatigue dulls outrage; enforcement lags become priced-in.

Step 4: Reinforcement. Money rewards firms that master the workaround; competitors must adopt or lose. Compliance becomes choreography; risk becomes actuarial.

Design response: break the loop upstream.

- **Reduce payoff (P):** election design (top-four + RCV) and small-donor financing blunt the marginal impact of any single funder's threat.
- **Increase cost (C):** exclusive-vendor rules, provenance tags, fast IE APIs, and platform archives raise legal and reputational risk.
- **Shorten detection lag (Δt):** real-time, machine-readable filings; public ad libraries with asset IDs; auto-diff of creatives and targeting.
- **Stiffen consequences (K):** presumptive-coordination standards, vendor disqualification, treble civil penalties, platform takedown duties.

Simple payoff model (intuition aid).
Let **EV = p·V − (C + K)** where *p* is probability of decisive influence, *V* electoral value, *C* operational/forensic cost, *K* expected legal/reputational penalty.
Policy aim: push **EV ≤ 0** for covert coordination by (i) lowering *p* (plurality-diluting reforms), (ii) raising *C* (instrumentation/provenance), (iii) raising *K* and shrinking detection lag.

Healthy norm reset. When coordination is **observable by default** and **punished on exposure**, professional incentives flip: the cheapest, safest strategy becomes staying truly independent—or competing on persuasion the hard way.

Heuristic for policymakers.
If a reform **only** tightens words (definitions) without changing **infrastructure** (assets, vendors, timing, archives), expect rapid workaround and accelerated norm drift. Build rules that operate where evasion must leave a trail.

4. Evidence Stack I: Adversarial Forensics

4.1 Creative-hash similarity for TV/digital assets

Goal. Detect "independent" creatives that are materially derived from campaign assets.

Ingest.

- Capture all aired/served assets: TV spots, CTV, YT, programmatic videos, display, radio, mail PDFs.
- Normalize: strip letterboxing, pad to common fps/size; extract audio as mono 16 kHz; OCR lower-thirds.

Hashes (use multiple, then fuse).

- Video frames: perceptual hashes per scene cut (pHash/dHash/aHash) + SSIM on keyframes.
- Audio: chroma-based fingerprints (e.g., CQT shingles), 2-D locality-sensitive hashing.
- Graphics: ORB/SIFT keypoints on slates, logos, B-roll; template hashes for color bars, slugs.
- Text: MinHash on OCR transcripts and supers; separate layer for disclaimers/slates.

Lineage graph.

- Nodes: assets; Edges: similarity with weights w_v (video), w_a (audio), w_t (text).

- Score S = alpha*w_v + beta*w_a + gamma*w_t, tuned on labeled sets.
- Flag if S ≥ tau and earliest-seen time delta ≤ 72 h across entities.

Operational rules of thumb.

- Tight windows: 24–72 h between campaign publication and IE airing raises suspicion.
- "Creative twins": identical VO cadence + lower-third template reuse + shared B-roll > any single cue.
- Versioning tell: same footage, different crop/aspect and color grade → deliberate obfuscation pattern.

False-positive control.

- Whitelist truly stock B-roll and public domain footage via provider IDs.
- Require at least two modalities to exceed thresholds (e.g., audio + lower-thirds).
- Human audit on borderline S in primaries within 14 days of the vote.

4.2 Impression-timing vs. floor action and markups

Goal. Show functional coordination by coupling IE delivery to legislative timing.

Event model.

- Define atomic events: bill introduction, committee markup start/stop, amendment posting, whip count leaks, floor scheduling, key procedural votes.
- Pre-register windows: W_pre = [−72 h, −6 h], W_vote = [−6 h, +6 h], W_post = [+6 h, +72 h].

Data.

- Impression logs by DMA/ZIP (TV, CTV, digital); spend timestamps; creatives IDs.
- Legislative calendars with minute resolution; whip reports where obtainable.

Tests.

- Burst detection: Poisson process baselines per DMA; flag λ spikes in W_vote against W_pre.
- Coupling ratio R = impressions(W_vote) / max(impressions(W_pre), epsilon).

- Permutation test: shuffle legislative timestamps within the week; compute p-value of observed R.
- Cross-correlation: impressions vs. event series; peak near lag 0 with CI.

Attribution safeguards.

- Require multiple alignments across distinct events or districts to claim pattern class.
- Control for major exogenous news shocks (national address, scandal days).
- Use pre-registered hypotheses to avoid p-hacking in high-volume cycles.

Interpretation tiers.

- Tier A (strong): $R \geq 3$ on ≥ 2 legislative events, $p \leq 0.01$, same IE + same creative family.
- Tier B (moderate): $R \geq 2$ once, corroborated by creative lineage or vendor overlap.
- Tier C (weak): timing alone; use as cue for deeper vendor/creative audit.

4.3 Vendor-lattice mapping and sub-vendor leakage

Goal. Expose "paper wall" separations where the production supply chain is shared.

Graph construction.

- Bipartite lattice $G = (E \cup V, \text{edges})$: entities E (campaigns, IEs, parties, 501(c)s); vendors V (prime shops, subs, freelancers, cloud workspaces, DSP seats).
- Enrich nodes with corporate family IDs (parent/affiliate), registered agents, domains, billing addresses, bank metadata (where lawful).

Signals.

- Common-agent hubs: vendors servicing both a campaign and an IE in the same race window.
- Sub-vendor bleed: distinct primes but identical post house, colorist, VO talent, or editing bay IDs.
- Cloud fingerprints: identical S3 bucket naming schemes, CDN paths, asset manifest formats, font packs.
- Temporal reuse: the same crew appears across entities within $T \leq 30$ days of each other's flight.

Metrics.

- Overlap coefficient $O(E1,E2) = |N(E1) \cap N(E2)| / \min(|N(E1)|, |N(E2)|)$.
- Family-aware overlap O* counting affiliates as one.
- Shortest path length $d(E_{campaign}, E_{IE})$ through vendor nodes; flag $d \leq 2$ within race window.
- Betweenness centrality of vendor nodes; high-B vendors are choke points.

Leakage typology.

- "Affiliate shuffle": one holding company owns both shops; different brand names, same back office.
- "Freelancer seam": a contractor straddles teams and carries creative files/templates.
- "Platform seat": shared DSP/ad-server seats with identical audience taxonomies or pixel IDs.

Controls & remedies.

- Exclusive-vendor rule proposal: a corporate family may not serve both sides of a race within a cycle.
- Sub-vendor disclosure: require primes to list subs and freelancers above 5% of the contract.
- Audit tags: platforms issue immutable seat IDs; broadcasters publish buy sheets with vendor IDs.

4.4 Message diffing: n-gram and script overlap windows

Goal. Quantify when IE messaging is functionally the same as campaign messaging in coordination-sized windows.

Corpora.

- Campaign: speeches, press releases, emails, stump scripts, web copy, organic social.
- IE: ad scripts, VO transcripts, mail copy, landing pages.
- Normalize to sentences; remove stopwords minimally; keep named entities.

Similarity stack.

- Lexical: TF-IDF cosine on 2- to 5-grams; Jaccard on shingles; MinHash LSH for fast pairing.
- Semantic: transformer embeddings with cosine similarity; sentence-level alignment.
- Stylometric: POS tag n-grams, cadence features (syllables per clause), VO pause histograms.

Windowing.

- Define W = 72 h rolling windows keyed to first appearance of a campaign phrase or frame.
- Compute sim scores S_lex, S_sem, S_style across W against IE outputs.
- Flag if S_sem ≥ 0.85 and S_lex ≥ 0.35 within W, or if stylometric cadence matches within 1 sd and VO actor class is identical.

Phrase lineage.

- Track rare phrase adoption (e.g., “tax on your groceries,” “crime of convenience”) via inverse document frequency across cycle corpus; rare-to-common jumps inside W are probative.
- Detect templated lower-thirds: identical slot grammar (“NAME — FAILED US ON [BILL]”) with placeholder substitution.

Decision rubric.

- Class 1: High-confidence mirror — multi-modal match (lexical + semantic + stylometric) inside 24–48 h, paired with creative lineage or vendor overlap.
- Class 2: Probable mirror — strong semantic match inside 72 h plus timing Tier B.
- Class 3: Parallel messaging — high similarity but outside 72 h or based on widely diffused national frames.

Minimizing overreach.

- Exclude generic policy boilerplate and national party talking points unless phrase rarity supports specificity.
- Require corroboration from at least one other layer (creative hashes, timing, vendor lattice) for public accusations.

Putting the stack together

- **Trigger.** Any one layer crossing a threshold opens a case file.
- **Corroborate.** Seek a second orthogonal layer within 48 h.
- **Score.** Compute a presumptive-coordination score: PCS = w1*Creative + w2*Timing + w3*Vendor + w4*Message.
- **Act.**
    - PCS ≥ tau_A → refer for enforcement, publish dashboard card with evidence slices.
    - tau_B ≤ PCS < tau_A → public yellow flag; request vendor and targeting attestations.
    - PCS < tau_B → retain as intel; no public action.

Ethic. The aim is not to criminalize persuasion but to make covert coordination operationally hard, quickly visible, and reputationally costly—so that the lowest-cost strategy is to stay truly independent.

5. Evidence Stack II: Causal Inference Without Confessions

5.1 Event studies around spending shocks

Purpose. Estimate the short-run causal effect of outside-money bursts on outcomes (polling, small-dollar inflows, sentiment, roll-call posture), without needing proof of coordination.

Design.

- Define shocks: first large IE flight, message switch, or DMA-level burst.
- Windows: pre = [−14d, −1d], impact = [0d, +7d], post = [+8d, +21d].
- Outcomes Y: district polls, ActBlue/WinRed small-donor counts, search trends, social sentiment, whip-sheet movements, ad recall.

Identification.

- Model: Y_it = a_i + b_t + sum_k beta_k * Shock_{i,t+k} + X_it' gamma + e_it
  where a_i = unit fixed effects, b_t = day/week fixed effects, X_it = controls (news, macro shocks).

- Parallel trends assumption: absent the shock, treated and comparison units move similarly.

Threats & fixes.

- Anticipation: pre-trends check; redefine shock as “first unexpected flight” using embargoed buy sheets.
- Spillovers: cluster by media market; drop adjacent DMAs in sensitivity.
- Seasonality/news: include national × day FE and topic-specific news dummies.

Robustness.

- Placebo shocks on pre-periods.
- Vary windows (±7d/±21d); restrict to nights/weekends.
- Outcome triangulation: if polling moves but donations/search do not, downgrade confidence.

Reporting.

- Pre-register shock definition, windows, and outcomes.
- Plot event-time coefficients beta_k with CIs; highlight pre-trend flatness and immediate impulse.

Interpretation.

- A sharp, specific jump in Y aligned with Shock and clean pre-trends → consistent with causal influence of the spend. Combine with Section 4 fingerprints for presumptive coordination.

5.2 Regression discontinuity in near-tie races

Purpose. Leverage knife-edge elections to estimate the causal effect of heavy IE support on representation or behavior.

Setup.

- Running variable r = vote margin (winner minus loser).
- Cutoff at r = 0 (just wins vs. just loses).
- Treatment: high IE support for candidate A in cycle t.

- Outcomes: subsequent roll-call scores, committee behavior, earmark success, donor mix, policy stances.

Identification.

- Local randomization: near the cutoff, assignment approximates random; units just above/below are comparable.
- Estimator: local linear RD with triangular kernel.
- Spec: Y = a + tau * 1[r >= 0] + b1 * r + b2 * r*1[r >= 0] + e within bandwidth h chosen by MSE/min-bias rules.

Key contrasts.

- Standard RD (incumbency effect): does winning per se change Y?
- Augmented RD: instrument subsequent behavior with an indicator for "win + high-IE support," or run a difference-in-RD comparing high-IE vs. low-IE close races.

Threats & fixes.

- Sorting/manipulation: McCrary density test around r = 0; audit precinct anomalies.
- Heterogeneous bandwidth sensitivity: report tau across multiple h; include robust bias-corrected CIs.
- Contamination: exclude recounts/irregularities; check balance on covariates within h.

Extensions.

- RD in time: close primaries as cutoff for access to national money.
- Fuzzy RD: if "treatment" (e.g., IE flood) rises sharply but not perfectly at the cutoff, use the cutoff as an instrument.

Interpretation.

- If "win with IE flood" vs. "lose with IE flood" shows durable shifts in Y (e.g., whip behavior, donor composition), attribute a causal effect to the IE-enabled victory channel, not merely candidate ideology.

5.3 Synthetic controls for policy flips

Purpose. Attribute post-election policy shifts (votes, sponsorships, committee priorities) to an IE-driven intervention in a specific district or member.

Design.

- Treated unit: member/seat experiencing a documented IE barrage and a subsequent policy flip (e.g., reverses stance on bill X).
- Donor pool: similar districts/members without the barrage in the same cycle.
- Pre-period: multiple terms or months to learn weights; Post-period: from IE onset to outcome stabilization.

Method.

- Build a weighted "synthetic twin" from donor pool that matches pre-treatment trajectories of key predictors (past votes on adjacent issues, ideology scores, district demographics, fundraising mix, media exposure).
- Estimate effect = Y_treated,post − Y_synth,post.

Assumptions.

- No unobserved shocks uniquely hitting treated unit post-intervention other than the IE.
- Good pre-period fit implies credible counterfactual.

Diagnostics.

- Pre-fit RMSPE; require low gap before treatment.
- Placebo tests: reassign "treatment" to donor units; the treated unit's post-gap should stand out (rank test).
- Sensitivity: leave-one-out on donor pool; keep effect sign/magnitude stable.

Data to include.

- Granular IE timelines, creative type (positive/negative), targeting reach.
- Policy outcome granularity: vote timing, co-sponsor date, whip notes.
- Covariates: district shocks (plant closures, disasters), leadership pressure indices.

Variants.

- Augmented synthetic control (ASC): combine outcome regression with synthetic weights for better small-N performance.
- Matrix completion: when outcomes are panel-structured with missingness.

Interpretation tiers.

- Strong: tight pre-fit, large and persistent post-gap, top-ranked vs. placebos, aligned temporally with IE onset.
- Moderate: decent pre-fit, medium gap, some placebo overlap but corroborated by Section 4 fingerprints.
- Weak: noisy pre-fit or short post window—treat as suggestive, not decisive.

Cross-method guidance (for 5.1–5.3)

- **Pre-register** designs to avoid researcher degrees of freedom.
- **Triangulate**: require at least two independent methods pointing in the same direction before strong claims.
- **DAG first**: write the causal graph; list back-doors you close and those you cannot.
- **Transparency**: release code, data schemas, and simulated stress tests (placebos, bandwidths, alternative outcomes).
- **Communication**: convert effects into intuitive units (e.g., "IE burst likely swung 1.5–2.0 pts in week t" or "raised probability of a policy flip by ~20%").

6. Chokepoint Design

6.1 Exclusive-vendor/affiliate rules (covering subs)

Purpose. Eliminate the "paper wall" where one corporate family quietly serves both a campaign and an IE in the same race.

Rule (model language).
"No vendor, subsidiary, affiliate, common-control entity, or retained subcontractor may provide strategic, creative, data, media buying, or production services to (i) a candidate committee and (ii) an independent-expenditure entity influencing the same race within the same election cycle."

Scope & definitions.

- *Affiliate/common control:* ≥20% ownership, shared officers/board, or shared back-office systems (finance, HR, cloud tenancy).

- *Subcontractor:* any entity delivering ≥5% of contract value or handling creatives/data.
- *Covered services:* strategy, polling, targeting, creative, editing, VO, media buying, ad ops, analytics, data engineering.

Compliance mechanics.

- Vendor registration with **Corporate Family ID**; mandatory listing of subs/freelancers.
- **Conflict API:** campaigns/IEs query vendor IDs before engagement; system returns "clear/blocked."
- **Attestation on invoice:** "No services to an IE/campaign in this race this cycle by any affiliate/sub."

Enforcement & penalties.

- Strict-liability civil penalties; disgorgement of contract; race-cycle disqualification for the vendor family.
- For willful breaches: treble damages + referral for criminal fraud if filings were false.

Edge cases & safe harbors.

- *Commodity services* (printing, postage) allowed if no access to creatives/targeting.
- Firewalled national vendors can serve **different races**; not both sides of one race.

Audit signals.

- Corporate-family graphs; shared registered agents; cloud/asset forensics; repeated talent overlaps.

6.2 Cooling-off windows for staff transitions

Purpose. Stop the "talent carousel" that transfers tacit strategy from a campaign to an IE (or vice-versa).

Rule (model language).
"No individual who has served a candidate, party, or coordinated committee in a covered race may provide services to an IE influencing that race (or its opponent) for **180 days** after their last access to nonpublic strategic materials. The same applies in reverse."

What counts as access?

- Nonpublic polling, messaging memos, creative drafts, targeting schemas, media calendars, vendor briefs, whip counts.

Mechanics.

- **Exit certification**: last day + list of materials accessed; auto-calculates cooldown end date.
- **Employment registry**: hashed resumes (role, employer, race, dates) to allow auditors to verify movements.
- **Role carve-outs**: compliance or purely administrative roles may cool off at **90 days** with additional attestations.

Enforcement.

- If a hire violates the window: cease-and-desist order; voided invoices; IE spend treated as coordinated (in-kind).
- Personal liability for falsified attestations; vendor disqualification next cycle.

Mitigations for small markets.

- Allow *non-overlapping portfolios* (e.g., school board vs. congressional) during cooldown.
- Subsidize talent pipelines to reduce hardship in rural states.

Audit signals.

- LinkedIn/Git/Slack timestamp drift; identical targeting taxonomies reappearing post-move; VO/creative fingerprints persisting.

6.3 Creative-provenance tags and broadcaster/platform duties

Purpose. Make asset lineage and delivery auditable in near real time.

Provenance standard.

- Every ad asset (TV/digital/audio/mail PDF) carries an embedded **Ad-Provenance Bundle (APB)**:
  asset_id • campaign/IE id • vendor id • corporate_family_id • upload_ts • checksum • creative_version • targeting_schema_id (digital) • disclaimer text.
- Cryptographically signed by the uploader; checksummed media stored by host.

Broadcaster/platform duties.

1. **Verify & store** APB on upload (reject if missing/invalid).
2. **Publish within 24h** to a public **Ad Library API**: asset_id, buyer, spend, flight, targeting slices (age/geo/context), vendor ID, corporate family ID, creative thumbnails/transcripts.
3. **Emit diffs**: if any field changes (e.g., retargeted audience), append a new version; never overwrite.
4. **Deliver logs**: impression-level aggregates by hour/DMA/zip for oversight (privacy-safe).
5. **Retain** assets/metadata ≥ 4 years.

Auto-flagging.

- The library runs **hash-similarity** and **n-gram** checks across entities; flags "creative twins" and "targeting twins."
- Vendor-overlap alerts when the same corporate family appears across a race's campaign and IE.

Compliance incentives.

- Noncompliant platforms face escalating fines and liability shields removed for disputed ads.
- Broadcasters receive **safe harbor** if they validate APB and publish on time.

Privacy & speech safeguards.

- Targeting fields disclosed at **category level** (age brackets, interests, context), not user-level data.
- No moderation of content; duty is **transparency**, not editorial control.
- Clear appeal pathway for mistaken flags; publish audit trails.

Interoperability.

- Open schema (JSON + checksums); shared **Vendor ID Registry** with corporate-family graph.
- State and federal regulators, journalists, and watchdogs consume the same API.

Success metrics.

- Median disclosure latency < 24h; % of ads with valid APB > 99%; reduction in unresolved "who paid/targeted?" cases; rise in detected vendor/creative overlaps **before** Election Day.

Design principle across 6.1–6.3: don't police intent—**constrain infrastructure.** When vendors can't straddle, staff can't hop fast, and every ad carries its pedigree in public, covert coordination becomes slow, risky, and expensive—the honest path becomes the cheapest path.

7. Move the Battlefield

7.1 Top-four + ranked-choice to dilute single-issue gatekeeping

Problem. Closed or low-turnout primaries let a concentrated funder bloc "discipline" incumbents with targeted IEs.

Intervention. Replace party primaries with a **single nonpartisan top-four primary** advancing to a **ranked-choice (RCV) general**.

Mechanism of relief.

- **Vote-splitting removed:** adjacent candidates can coexist without cannibalizing each other; gatekeepers can't "pick the spoiler."
- **Broader median:** a wider electorate selects finalists; single-issue wedges have less pivot power.
- **Downside protection:** candidates risk less by defying a funder because second/third-choice flows can offset a single bloc's punishment.

Implementation notes.

- Ballot design: bubble-fill rank columns; overvote prevention copy; exhausted-ballot education.
- Count pipeline: publish **round files** (CSV) with precinct-level transfers; open-source tabulation code; risk-limiting audits adapted to RCV.
- Transition period: dual runs (plurality + simulated RCV) for 1 cycle to build trust; publish delta analyses.

Metrics.

- Share of finalists with cross-coalition support.
- Reduction in IE concentration in primaries (HHI index).
- Elasticity of vote share to last-minute negative IE spend (pre vs post reform).

Risks & mitigations.

- Spoiled ballots → invest in voter education + clear instructions.
- Party resistance → guarantee party labels on ballots; preserve party endorsement lines.

7.2 Democracy vouchers / small-donor matching as counter-leverage

Problem. The credible *threat of negative spend* shapes votes before a dollar moves.

Intervention A — Vouchers.

- Each eligible voter receives **voucher credits** (e.g., $25×4) assignable to qualifying candidates.
- Candidates qualify by thresholds (signatures + small-donor seed).

Intervention B — Matching.

- Publicly match in-district donations ≤ $200 at **6:1–10:1**, with caps and instant verification.

Mechanism of relief.

- **Pre-capitalization:** challengers and incumbents signal *bankable* grassroots support; the marginal effect of a hostile IE drops.
- **Counter-threat:** organized communities can rapidly **redeem vouchers** or trigger matching bursts in response to punishment campaigns.
- **Network inversion:** time with constituents dominates time with bundlers; "access" shifts.

Design details.

- Real-time wallet/API: voters allocate on web/mail; candidates see **available but unspent** voucher pools (privacy-safe).
- Anti-fraud: KYC light (address + last-4), audit flags for rapid, correlated redemption.

- Equity guardrails: residency weighting; prohibit conversion to IEs; publish voucher flows daily.

Metrics.

- % of total receipts from vouchers/matched smalls.
- Number of viable non-incumbent entrants.
- Response time from negative IE to grassroots counter-funding.

Risks & mitigations.

- "Astro-grass" laundering → strict in-district rules; random audits; clawbacks.
- Program capture → per-cycle caps; sunset + reauthorization contingent on evaluation.

7.3 Fast-disclosure infrastructure: 24-hour IE & ad APIs

Problem. By the time spending is visible, the election has moved; covert coordination enjoys a **first-mover window**.

Intervention. Statutory **T+24** disclosure for IEs and paid political communications, backed by **machine-readable APIs**.

What must be filed (within 24 hours of first run/serve).

- Buyer (entity, treasurer), amount committed, vendor family ID.
- Asset IDs (with cryptographic checksums), first air/serve time, markets/targets (category-level).
- Funding source since last report (aggregated donors if super-PAC; top 10 listed by amount).

API/registry requirements.

- **Ad Library API** (read): creatives, transcripts, thumbnails, spend by hour/DMA/ZIP, targeting categories, provenance bundle (see §6.3).
- **IE Filing API** (write): JSON schema with validation; immediate public mirroring; versioning on amendments.
- **Vendor Registry:** canonical IDs with affiliate graph; open query.

Why speed matters.

- Shrinks detection lag Δt → lets media, watchdogs, and opponents surface vendor overlaps, creative twins, and timing couplings **before** consequential votes.
- Converts covert tactics into **live reputational costs** (press, dashboards, debates).

Operational supports.

- Bulk upload & CLI tools for filers; sandbox + conformance tests.
- Auto-flaggers: hash similarity, vendor overlaps, timing spikes; public webhooks for alerts.
- Accessibility: Open data license; daily CSV dumps for non-technical users.

Metrics.

- Median latency from first impression to public visibility (< 24h).
- % filings passing validation on first submission (> 95%).
- Number of investigative pieces/dashboards citing API within cycle.
- Share of flagged overlaps **pre-election** vs **post-election**.

Risks & mitigations.

- Filing burden → provide free tooling; safe-harbor for good-faith filers who fix within 48h.
- Over-flagging → tiered alerts; require multi-modal corroboration (hash + vendor OR timing).
- Platform noncompliance → escalating fines; conditioning carriage on real-time feeds for broadcasters.

Throughline. Changing *who advances*, *who can be financed quickly*, and *how fast the public sees the money and the message* lowers the marginal power of single-issue gatekeepers and short-circuits stealth coordination. The battlefield moves from back-rooms and time lags to open APIs and voter-centered finance.

8. Litigation & Enforcement Wedges

8.1 Firewall fraud and deceptive compliance

Problem. "Independent" spend is laundered through paper firewalls: different LLCs, letterheads, and memos—same hands, assets, or cloud stack.

Theory of the case.

- **Deceptive practices / false filings:** When entities certify independence while sharing creatives, staff, or data, filings become materially misleading.
- **Constructive coordination:** Functional integration (shared production chain, mirrored assets within 24–72 h) defeats independence even absent a smoking-gun email.

Model provisions (civil).

- **Presumptive coordination:** If an IE and a campaign in the same race share (i) a vendor corporate family or sub-vendor, (ii) substantially similar creative assets (hash lineage), and (iii) synchronized flighting around official acts, the IE is *presumptively coordinated* unless rebutted by documentary proof of separation (asset logs, access controls, provenance).
- **False attestation liability:** Any officer signing independence attestations is strictly liable for material misstatements; knowing violations trigger treble damages.

Discovery targets.

- Build manifests, project files, font packs, LUT/color profiles, version trees, cloud bucket paths, export metadata, edit-bay calendars.
- HR/onboarding records, device logs, contractor invoices, per-seat ad-platform IDs.

Remedies.

- **In-kind reclassification:** Convert the IE to a contribution; if over limits or prohibited source → disgorgement + civil penalties.
- **Equitable relief:** Emergency injunction halting flights; vendor disqualification for cycle; mandatory corrective disclosure.

Safe harbors & due process.

- Commodity vendors (mailhouses, printers) exempt if no access to targeting/creative.
- Right to submit counter-forensics (independent build logs, third-party audits).

- Appeals window with public docket to prevent weaponized accusations.

Metrics.

- Rate of pre-election injunctions; % of presumptions rebutted; cycle-to-cycle drop in shared-vendor patterns.

8.2 Conduit schemes (LLCs, straw donors): bank/wire metadata

Problem. True sources are disguised via reimbursed donors, pop-up LLCs, and pass-through nonprofits.

Theory of the case.

- **Contribution in the name of another:** Classic straw-donor violation.
- **Wire fraud / money laundering hooks:** False representations in filings and structured transfers to conceal origin.

Detection playbook.

- **Entity forensics:** Secretary-of-State pulls; registered agent reuse; simultaneous formation clusters; no revenue history.
- **Funds-flow graphs:** Banking metadata (SWIFT/ACH), wire timing, memo fields; circularity (A→B→C→A), burst patterns pre-vote.
- **Means-test screens:** Donor incomes vs. contribution levels; correlated reimbursements (same-day ATM/deposits).
- **Common-paymaster clues:** One law firm or consultant paying many “donors” within 48 h.

Model provisions (criminal/civil).

- **Strict disgorgement:** Contributions traced to conduits are void; recipients disgorge to the state.
- **Multiplier penalties:** 2×–3× the routed amount on knowing schemes; eligibility bans for donors/entities for two cycles.
- **Officer liability:** Treasurers and controllers who “should have known” face civil penalties for willful blindness.

Discovery & subpoenas.

- Bank records (KYC, wires, checks, deposit slips), retainer agreements, invoice trails, payment processors' merchant data.
- Cross-match with vendor invoices to spot "refund" loops.

Whistleblower levers.

- Anonymous portals + bounties tied to recovered amounts; safe harbor for first cooperator.

Defenses & guardrails.

- Genuine LLC donors allowed if operating businesses with normal financials and board approvals.
- Charity pass-throughs prohibited only when earmarking exists; otherwise disclose as "source-obscured" with heightened audit (policy choice).

Metrics.

- Time-to-trace for suspect flows; % of pop-up LLCs sanctioned; recovered dollars vs. enforcement cost.

8.3 State-level transparency and archive mandates

Problem. Federal pace is slow; crucial data (creative assets, targeting, spend timing) isn't reliably public in time to matter.

State tools (fast and bitey).

- **Ad Library Mandate:** Any political ad reaching ≥1,000 state residents in 24 h must be filed with the state's **Public Ad Library** within **T+24**: asset + transcript + spend + vendor family + targeting categories + funding since last report.
- **Platform Duties:** Broadcasters, CTV, social platforms must validate **Ad-Provenance Bundles** (see §6.3), publish via API, and retain for 4 years.
- **IE Rapid Reports:** IEs file commitments (not just disbursements) ≥ $10,000 within 24 h; amendments on change ≥ 10%.

Enforcement.

- **Per-day per-ad fines** for late/missing filings; escalating for repeat offenders; public "noncompliance board."
- **Carriage conditions:** State can condition cable/broadcast carriage on compliance (where authority exists).
- **Private right of action:** Verified residents and candidates can sue for injunctive relief + fees on noncompliant platforms/spenders.

Interoperability & access.

- **Open schema/API:** JSON endpoints for creatives, spends, targets; daily CSVs for journalists.
- **Vendor Registry:** state-maintained IDs with affiliate graphs; bulk diff feeds for watchdogs.
- **Webhooks & alerts:** Auto-notify on hash-similarity flags, vendor overlaps, timing spikes.

Safeguards.

- Privacy: disclose targeting at category/geo granularity (no user-level).
- Speech neutrality: transparency duty only; no content moderation mandates.
- Error correction: 48 h safe harbor to cure good-faith mistakes; visible edit history.

Model statute snippets.

- "Any person or platform disseminating a covered communication shall, within 24 hours of first dissemination, file a provenance bundle and make it available via the Public Ad Library API. Failure to do so constitutes an unfair and deceptive act in commerce."
- "An independent expenditure reported after T+24 shall be presumed strategically timed to evade scrutiny; such presumption may be rebutted only by clear and convincing evidence."

Metrics.

- Median latency (impression→public file); % of ads with valid provenance; number of pre-election investigations using state APIs; decline in unresolved "who paid?" stories.

Strategic throughline. Aim litigation and mandates where evasion must leave traces: **assets, wires, vendors, timestamps.** Treat false firewalls and conduits as *fraud problems* with documentary remedies; use state transparency to collapse the time it takes for truth to surface.

9. Whistleblowers Who Actually Matter

9.1 Protecting junior staff and vendors

Why juniors matter. Real coordination lives in edit bays, DSP seats, cloud folders, and invoice trails—places staffed by junior producers, ad ops specialists, media traffickers, and freelance subs. They see the practical breach (asset reuse, "separate" logins on the same box), but they're the least protected.

Protection architecture.

- **Role-based immunity:** Statute/regulation that shields *individual contributors* (ICs) who report good-faith evidence of coordination/conduit schemes from civil liability tied to NDAs, work-product claims, or non-disparagement clauses.
- **Anti-retaliation with teeth:** Automatic presumption of retaliation if adverse action occurs within 180 days of a protected disclosure; remedies include double back pay, reinstatement priority, and vendor debarment for offending employers.
- **NDA safe clause (model line):** "Nothing in this agreement prohibits disclosure of suspected election-law violations to regulators, law enforcement, or designated transparency portals; copies of creative assets and logs necessary to substantiate such disclosures are authorized."
- **Vulnerable-vendor shield:** Extend protections to sub-vendors/freelancers (post houses, VO talent, colorists, copy editors) whose contracts are often at-will; blacklisting becomes an actionable tort under the program.

Operational supports.

- **Evidence handling kit:** Plain-language guidance on what to capture legally: asset checksums, export metadata, project file version trees, seat IDs, timecards, buy sheets, and redacted invoices (no client PII beyond what is necessary).
- **Escrowed counsel:** State-authorized panel of pro bono attorneys on retainer; triage within 48 h; privileged intake.
- **Protected device pathway:** Upload tools that strip personal identifiers, compute hashes locally, and generate a tamper-evident manifest before submission.

Guardrails (to prevent weaponization).

- Perjury-backed attestations for submitters.
- Penalties for knowingly false claims; clear appeal path for accused firms.
- "Need-to-know" redaction defaults, with full materials accessible only to regulators/courts.

Success metrics.

- Median time from disclosure to regulator intake decision (< 10 days).
- % of cases with corroborating technical artifacts (hashes, logs).
- Retaliation rate year-over-year; number of reinstatements and vendor debarments.

9.2 Bounties, safe channels, audit-ready pledges

Bounties (make truth liquid).

- **Structure:** Scaled awards (e.g., 10–30% of penalties/disgorgement) for information that results in confirmed findings of illegal coordination, conduit donations, or deceptive filings.
- **Fast-track micro-bounties:** Fixed $1–5k payments for verifiable, discrete artifacts (e.g., a matching creative lineage with source files, a seat-ID screenshot proving shared platform access) that trigger formal inquiries.
- **First-cooperator advantage:** Larger shares for the first whistleblower whose evidence materially advances the case; subsequent corroborators get smaller, guaranteed stipends.

Safe channels (make it easy and safe to talk).

- **Dual rails:** (1) Regulator portal with legal privilege; (2) Independent public-interest trustee (university/NGO) that time-stamps and escrows evidence, forwarding only with the discloser's consent or subpoena.
- **Anon-by-default intake:** Secure uploads (Tor/HTTPS), stateless forms, no phone/email required to file an initial cryptographic receipt; optional contact enabled via one-time relays.
- **Chain-of-custody receipts:** Automated manifests (file hashes, timestamps, signer keys) returned to the whistleblower and stored with the trustee; admissibility relay for courts.

Audit-ready pledges (flip the default for candidates and IEs).

- **Public pledge:** Candidates, parties, IEs, and vendors publish an *Audit-Ready Attestation* committing to:
    1. Maintain immutable creative and export logs;
    2. Publish vendor/sub-vendor lists and corporate-family IDs;
    3. Disclose platform seat IDs used for buys;
    4. Cooperate with provenance checks within 72 h of a credible flag;
    5. Forbid retaliation; include the NDA safe clause in all contracts.
- **Verification badge:** Independent auditors issue time-bound badges; loss or refusal is itself a public signal.
- **Red-button protocol:** If a staffer invokes retaliation protections, the pledging entity must freeze related buys/creatives for 72 h and produce a limited log set to the auditor/regulator.

Incentives & enforcement.

- **Procurement carrots:** Broadcasters/platforms prioritize inventory access or fee discounts for pledged, badge-holding buyers.
- **Penalty multipliers:** Treble penalties when violations occur *after* a signed pledge; bounty shares increase correspondingly.
- **Platform hooks:** Carriage/platform policies require provenance compliance and honor red-button freezes upon verified alerts.

Public feedback loop.

- **Disclosure scorecards:** Dashboards track who has active pledges, audit results, and response times to flags.
- **Case digests:** Anonymized quarterly summaries of whistleblower outcomes, bounties paid, and systemic fixes adopted.

Ethic and aim.
Whistleblowing should not require martyrdom. By pairing **legal shelter** for juniors and vendors with **real money**, **trusted, low-friction channels**, and **audit-ready commitments** from campaigns and IEs, we convert private knowledge into public accountability—fast enough to matter *before* the election is decided.

10. Public Dashboards & Narrative Tools

10.1 Pay-to-Vote Index (outside-money shocks vs. roll-call flips)

Purpose. A single, interpretable score that indicates whether recent outside-money shocks plausibly shifted a member's vote posture.

Inputs (per member m, bill b, window W).

- IE_shock: standardized spend/impressions change in W (e.g., z-score by DMA and cycle).
- Msg_align: semantic alignment between IE creative and the bill's salient frames (0–1; from §4.4).
- Timing_couple: coupling ratio of impressions to markup/floor windows (from §4.2).
- Baseline_prob: predicted "aye/nae" probability absent shock (whip/history model).
- ΔPosition: observed change in public position (no/lean/undecided/yes) or whip sheet updates.
- Counter_spend: standardized supportive spend from the other side in W.

Index (bounded 0–100).

1. Construct a raw effect score
   $R = w_1 \cdot IE_shock + w_2 \cdot Msg_align + w_3 \cdot Timing_couple - w_4 \cdot Counter_spend$
2. Convert to *swing likelihood* via logistic with baseline:
   $\hat{p} = \sigma(a \cdot R + b \cdot Baseline_prob)$
3. Map to index and penalize inconsistency:
   $PVI = 100 \cdot \hat{p} - \lambda \cdot |\Delta Position_expected - \Delta Position_observed|$

Tiering (display logic).

- 70–100: **High**—outside money very likely altered posture; show evidence slices.
- 40–69: **Medium**—concerning but needs corroboration (vendor/creative).
- 0–39: **Low**—no actionable coupling.

UI notes.

- Sparkline of IE spend vs. markup/floor timestamps.
- "Why this score?" explainer chips: IE_shock ↑, Timing_couple ↑, Counter_spend ↓.
- Downloadable methods card (YAML): weights, windows, last retrain.

Guardrails.

- Pre-register weights and windows per chamber/session.
- Confidence band rendered when inputs are missing/partial.
- Prominent disclaimer: *normative signal, not legal finding*.

10.2 Race dashboards: ad-hash overlap, vendor-overlap meters, 48-hour IE tickers

Purpose. Live, race-level situational awareness that surfaces presumptive coordination *before* ballots are cast.

Core tiles.

1. **Ad-Hash Overlap Meter**
    - Gauge: 0–100 based on fused audio/video/text similarity across entities (from §4.1).
    - Tooltip: "Matched VO cadence + lower-third template + B-roll lineage detected."
    - Click-through: side-by-side creative diff and time delta.
2. **Vendor-Lattice Map**
    - Bimodal graph (entities ↔ vendors) with affiliate highlighting.
    - Meter: overlap coefficient O* and shortest path d ≤ 2 flagged in orange/red.
    - Toggle: show subs/freelancers and corporate family roll-ups.
3. **48-Hour IE Ticker**
    - Rolling feed of IE filings (T+24 standard) with commit amount, asset IDs, targets.
    - Heat band shows bursts aligned to legislative or news events.
4. **Timing Coupler**
    - Cross-correlation plot: impressions vs. event series; peak near lag 0 is annotated.
5. **Messaging Twin Finder**
    - Phrase lineage panel with rare-phrase adoption curves and n-gram matches (from §4.4).

Alert logic (public-facing).

- **Yellow flag:** Any two layers cross moderate thresholds (e.g., Ad-Hash ≥ τ_1 and d ≤ 2).
- **Red flag:** Three layers cross strong thresholds within 72 h.
- Flags open a case card with evidence slices, timestamps, and rebuttal upload link.

Data hygiene & privacy.

- Display targeting at **category/geo bins** only.
- Obfuscate staffer names; show roles and dates.
- Keep raw creative assets/transcripts downloadable for verification.

Comparability.

- Per-race percentile ranks for each meter.
- Cycle-to-cycle deltas to show whether a jurisdiction is improving post-reform.

Provenance.

- Every tile shows data source (platform API, broadcaster file, IE filing) and hash of the ingested record.

10.3 Translating forensics into citizen-readable signals

Narrative principle. People act on *stories with receipts*. Compress technical forensics into plain, fair signals with verifiable slices.

Design patterns.

- **Traffic-light labels:** Green (no overlap), Yellow (investigate), Red (presumptive coordination).
- **Receipts first:** Each label expands to three screenshots: (1) creative diff, (2) vendor graph snippet, (3) timing plot with event markers.
- **Two-sentence summaries:**
    - "In the last 48 hours, an outside group aired ads that reuse the campaign's footage and captions, and used the same post house. The flights peaked during a committee markup on the same bill."

Citizen tools.

- **Shareable "race cards":** One-page PNG/PDF with QR to the live dashboard; optimized for social and local newsrooms.
- **Explain-this buttons:** Micro-glossary popovers ("What is vendor overlap?", "What counts as coordination?").
- **Confidence ribbons:** Low/Medium/High with a hover panel showing which layers fired and which did not.

Media & watchdog integrations.

- **Embeddable widgets:** iFrame/JS components for local outlets; static image fallbacks for print.
- **Webhook alerts:** Opt-in notifications to journalists/citizens when a race flips from Green→Yellow or Yellow→Red.
- **Case notebooks:** Auto-generated, time-ordered dossiers (Markdown/HTML) that compile plots, filings, and hashes—ready for FOIAs or complaints.

Tone & fairness.

- Avoid accusatory verbs; use *"signals suggest," "overlaps indicate," "presumptive under policy X."*
- Equal-opportunity lens: show counterspeech and corrections prominently; log retractions with timestamps.

Accessibility.

- Color-blind safe palettes; text alternatives for charts; mobile-first layout.
- Read-level target: 9th–10th grade for summaries; technical annexes one click away.

Impact metrics.

- Share of alerts issued **before** key votes or early voting.
- % of red flags that lead to rebuttal uploads, agency inquiries, or spend pullbacks.
- Dwell time and completion rate on "Why this score?" panels (proxy for comprehension).
- Cycle-over-cycle decline in high-risk overlaps (goal: shift behavior, not just expose it).

Ethic. The dashboards should *price the workaround in public*. When citizens can *see* the overlap, *verify* the receipts, and *understand* the stakes in two sentences, covert coordination loses its best weapon: time.

11. Ethics of Design: Making the Honest Path Cheap

11.1 Minimizing false positives and chilling effects

Principle. Our goal is deterrence-by-clarity, not paralysis-by-fear. Systems must flag *patterns*, not punish proximity.

Design guardrails

- **Two-layer rule.** Public flags require at least **two orthogonal signals** (e.g., creative lineage + vendor overlap, or timing coupling + messaging twin). Single-signal hits stay internal for monitoring.
- **Window discipline.** Use narrow, pre-registered windows (e.g., 24–72 h around events) to avoid "similar because topical" traps.
- **Whitelist commons.** Stock B-roll, public-domain footage, and national boilerplate don't count toward creative similarity scores.
- **Rarity weighting.** Emphasize rare phrases/visual motifs; de-weight ubiquitous frames ("cut inflation," waving flag).
- **Human-in-the-loop.** Borderline cases (within 10% of thresholds) require analyst sign-off before any public label.
- **Confidence ribbons.** Every public card carries Low/Medium/High confidence and states *which* layers fired and which did not.

Speech respect

- **Content neutrality.** Systems score *structure* (hashes, timing, vendors), not policy positions.
- **No gag effect.** Offer a safe "parallelism" lane: campaigns can pre-register shared public resources (issue toolkits) that reduce the chance of penalizing legitimate alignment.

Operational hygiene

- **Pre-registration.** Publish thresholds, windows, and weightings before the cycle; archive any updates with diffs.

- **Audit trails.** Immutable logs for data pulls, model versions, and analyst actions; external auditors can replay a flag from raw inputs.

Metrics to track chilling

- Ad volume and vendor diversity before/after reforms.
- Time-to-air for legitimate creatives (should not rise materially).
- Share of reversed flags and reasons (data error vs. judgment).

11.2 Due process for accused campaigns and PACs

Principle. Fairness sustains legitimacy. Accused entities get a fast lane to context, correction, and appeal.

Procedural guarantees

- **72-hour rebuttal window.** Red/yellow flags open a **secure rebuttal portal** where entities can upload counter-evidence: build logs, asset manifests, vendor attestations, staffing rosters, provenance bundles.
- **Presumption tiering.** Public language is conditional ("signals suggest," "presumptive under rule X") until rebuttal closes.
- **Independent review.** Rebuttals routed to an external reviewer pool (rotating, conflict-screened). Majority vote required to sustain a red flag.

Evidence standards

- **Document over assertion.** Time-stamped project files, edit metadata, cloud access logs, and platform seat IDs outweigh general denials.
- **Chain-of-custody checks.** Submitted materials get hashed and time-stamped; tampering triggers separate findings.

Outcomes & remedies

- **Sustain.** Flag remains; case dossier expands with rebuttal artifacts; referral to regulator if thresholds met.
- **Downgrade.** Red→Yellow or Yellow→Green with a visible rationale ("shared stock footage; no vendor linkage").

- **Retract & notify.** Prominent retraction badge; automated notifications to subscribers and media who embedded the widget.

Anti-weaponization

- **Complainant accountability.** Third-party complaints must include at least one verifiable artifact; repeated bad-faith filings lose portal privileges.
- **Rate limiting.** Cap public flags per race/day to prevent "alert flooding."

Right to context

- **Statement slot.** Accused parties may append a 150-word contextual statement to the public card during review.
- **Public provenance.** If a platform or broadcaster caused a metadata error, the dashboard attributes the fault upstream.

Equity considerations

- **Small-shop grace.** Micro-vendors get a guided template to produce acceptable logs without enterprise tooling.
- **Translation & accessibility.** Rebuttal portal supports multiple languages, screen readers, and mobile uploads.

Transparency after the fact

- **Case notebooks.** Publish a final, linkable notebook: raw inputs, methods, rebuttal artifacts, decision record.
- **Lessons learned.** Aggregate, anonymized post-cycle report on false positives, rule tweaks, and model retrains.

Success metrics

- Median time from flag to decision; fraction of flags sustained; reversal rate; stakeholder satisfaction surveys.
- Enforcement "hit quality": share of sustained reds that lead to voluntary changes (pulled ads, vendor swaps) or regulator action.

Ethic. Make cheating hard and honesty easy—*and* make fairness obvious. A system that can admit error, show its work, and dignify rebuttal is one that the public (and even the accused) can live with.

12. Case Typology (Abstracted)

12.1 “Clone creatives” pattern

Essence. Nominally independent ads reuse a campaign’s creative DNA—footage, lower-thirds, VO cadence, or scripts—within a short window.

Canonical signature

- **Multimodal overlap:** audio fingerprint + keyframe pHash + OCR’d supers match patterns (≥2 layers above threshold).
- **Tight deltas:** IE first-air ≤ 24–72 h after the campaign asset appears.
- **Version tells:** same B-roll with alternate crop, LUT, or aspect ratio (an obfuscation motif).

Likely pathway

- Shared production chain (post house, editor, VO talent) or deliberate “public cue” harvesting by a friendly shop.

Primary risks

- Illegal coordination if coupled with shared vendors or synchronized timing; reputational hit even if technically legal.

Defenses offered

- “Public materials safe harbor,” stock footage, national-party style guides.

Adjudication rubric

- **Red:** ≥2 creative layers match + vendor overlap or provenance logs missing.
- **Yellow:** strong matches but only public-source explanations; request build logs.
- **Green:** similarity limited to stock/party boilerplate; windows > 72 h.

Remedies

- In-kind reclassification; vendor disqualification for cycle; publish side-by-side diffs and APB gaps.

12.2 “Common agent” pattern

Essence. Separate entities claim firewalls but operational control or know-how flows through the same corporate family or intermediaries.

Canonical signature

- **Vendor lattice:** campaign ↔ vendor family ↔ IE path of length ≤ 2 during the race.
- **Sub-vendor bleed:** different primes share the same editor bay, colorist, VO booth, or cloud workspace.
- **Process twins:** identical audience taxonomies, file naming schemes, font packs, or export presets.

Likely pathway

- Holding-company agencies, franchise affiliates, or freelancers straddling teams; "paper walls" with shared back office.

Primary risks

- Functional coordination without explicit messages; false filings/attestations.

Defenses offered

- Firewall memos, separate LLCs, "no-contact" attestations.

Adjudication rubric

- **Red:** affiliate/vendor overlap + creative or messaging twin within 72 h.
- **Yellow:** affiliate overlap alone; demand attestations + technical audits.
- **Green:** commodity-only overlap (print/postage) with no access to strategy or assets.

Remedies

- Presumptive coordination unless rebutted with logs; disgorgement; treble penalties for false attestations; future-cycle conflict bans.

12.3 "Timed barrage" pattern

Essence. Outside spending is precisely synchronized to legislative moments (markups, procedural votes) to compel or reward member behavior.

Canonical signature

- **Coupling ratio R:** impressions in [−6 h, +6 h] around the event ≥ 2–3× baseline.
- **Repeatability:** same pattern across multiple events or districts; peak cross-correlation ≈ 0 lag.

- **Message align:** creative frames reference the bill's salient issue language.

Likely pathway

- Calendars and whip intel transmitted via public signals or intermediaries; IE holds creatives ready to surge.

Primary risks

- Compellence that skirts coordination on paper; can flip roll-call posture without explicit quid-pro-quo.

Defenses offered

- "We follow the news cycle," routine flight optimization.

Adjudication rubric

- **Red:** R ≥ 3 on ≥ 2 events + message align + any second layer (vendor/creative).
- **Yellow:** single-event spike with align; monitor for repetition.
- **Green:** broad, steady flights unrelated to legislative cadence.

Remedies

- Rapid disclosure violations trigger presumptions; injunctions on surging assets; dashboard red flag + Pay-to-Vote note.

12.4 "Conduit carousel" pattern

Essence. True sources hidden via straw donors, pop-up LLCs, or pass-through nonprofits; money cycles through a "carousel" before landing.

Canonical signature

- **Entity forensics:** same registered agent/address across many LLCs; formation bursts near filing deadlines.
- **Funds-flow loops:** A→B→C→A circular wires; same-day reimbursements; memo fields indicating "consulting" with no deliverables.
- **Means mismatch:** donors with low reported income making maxed contributions across multiple races.

Likely pathway

- Controller or law firm operating a hub-and-spoke scheme; reimbursements disguised as fees or expenses.

Primary risks

- Contribution-in-the-name-of-another; money laundering; false filings.

Defenses offered

- "Legitimate businesses," pooled fundraising, independent charitable grants.

Adjudication rubric

- **Red:** bank metadata confirms reimbursements or circularity + shell indicators.
- **Yellow:** shell indicators without bank proof; subpoena threshold met.
- **Green:** operating businesses with revenue, board approvals, and clean books.

Remedies

- Strict disgorgement from recipients; multiplier penalties; eligibility bans for donors/entities; treasurer liability; whistleblower bounties tied to recovered sums.

Using the typology

- **Intake:** classify incoming flags by pattern; pre-load expected evidence (hashes, vendor graph, wire data, timing plots).
- **Scoring:** compute a Presumptive Coordination Score (PCS) with pattern-specific weights (Clone > Common Agent > Timed Barrage > Conduit, context-dependent).
- **Action:** Red → referral + public dossier; Yellow → request logs/attestations + monitored card; Green → archive, no publicity.
- **Learning loop:** after each adjudication, update thresholds and whitelists to reduce future false positives.

13. The Lighthouse Protocol (Poetic Interlude)

13.1 Metaphor as design spec: beams, fog, currents, wakes

The **beam** is not a sermon; it is a clock.
It sweeps on schedule, numbered and public, so no captain can claim surprise.
In code: publish rounds, thresholds, and windows; keep the metronome audible.

The **fog** is policy—dense, shifting, plausible.
You do not argue with fog; you measure visibility and print it in the almanac.
In code: confidence ribbons, rarity weights, two-layer rule; show how far we can see today.

The **currents** are money—cold rivers under the hull.
Currents do not confess; they reveal themselves in drift and arrival times.
In code: T+24 filings, impression curves, coupling ratios; the tide chart that says where the pull begins.

The **wakes** are history—what a hull writes after it passes.
A wake is a signature: B-roll grains, VO cadence, vendor fingerprints.
In code: creative hashes, vendor lattices, provenance bundles; an index of the boat by the way it disturbs the water.

A lighthouse has **no favorites**.
It names beams, fog, currents, wakes—and lets the harbor do the rest.
Design spec: enumerate invariants, bind them to clocks, and let receipts sing.

13.2 Instrument the shoreline; don't sermonize the sea

Build **posts** before you write **poems**.

Nail a post at **timing**: every flight and filing gets a timestamp, a checksum, a public echo.
Set a post at **origin**: every creative carries its lineage in the open.
Drive a post at **contact**: vendors declare families and subs, staff cool their heels, seats wear serials.
Add a post at **redress**: a rebuttal door that opens fast, logs that replay the storm, badges that can be lost and won.

Between posts, string **lines** the public can see:

- A **line of sight**: dashboards that turn math into plain weather—green, yellow, red with receipts.
- A **line of custody**: hashes and manifests that outlast press cycles and memory.

- A **line of care**: shields for juniors and vendors, bounties that turn risk into livelihood.

Then step back.
Let ships argue policy in open water while the **shoreline instruments** price deceit:
make the cheap path the honest one, the clever workaround slow, noisy, and dear.

If you must speak, say only what the tide gauge says.
Leave adjectives to the wind.
Truth arrives as **clock time over water**: a sweep of light, a mark on a pole, a wake recorded and compared.

14. Implementation Roadmap

14.1 Data standards, IDs, upload flows, and public APIs

Core identifiers (stable, human-auditable, machine-safe)

- **entity_id**: canonical for campaigns, IEs, parties, vendors, subs, and platforms. Format: ENT-<namespace>-<base32-8>.
- **family_id**: corporate-family root (affiliates roll up here).
- **asset_id**: content-addressed adz:<sha256-16> of normalized creative.
- **seat_id**: ad-tech buying seat (DSP/ad-server) serialized by platform.
- **bundle_id**: provenance bundle digest apb:<sha256-16> of the JSON payload.
- **filing_id**: immutable IE filing record with version suffix vN.

Minimal open schema (APB v1, JSON)


```
{
  "bundle_schema": "apb_v1",
  "asset_id": "adz:4C9F…",
  "uploader_entity_id": "ENT-IE-7X2K9MFT",
  "vendor_id": "ENT-VND-5P8Q2ZJL",
  "family_id": "FAM-1R3T8",
  "media": {"kind": "tv|ctv|video|audio|display|mail_pdf", "duration_s": 30, "lang": "en"},
  "disclaimer": "Paid for by …",
```

"creative_version": "C-2026-AL-07-A",

"checksums": {"sha256": "…", "frame_phash": ["…"]},

"targets": {"geo": ["DMA:803","ZIP:02139"], "categories": ["age_25_54","context_news"]},

"first_air_ts": "2026-10-05T14:32:09Z",

"funding_since_last": [{"source": "donor_agg", "amount_usd": 250000}],

"sign": {"alg": "ed25519", "key_id": "KEY-ENT-IE-7X2K9MFT", "sig_b64": "…"}

}

Upload flows (T+24)

1. **Preflight**: publisher validates schema + signature; computes asset_id; de-duplicates.
2. **Accept**: return bundle_id, persist media + transcript, emit webhook.
3. **Publish**: API surfaces record within 24h; amendments create new versions (no overwrites).
4. **Linkage**: auto-attach broadcaster buy sheet (TV) or impression aggregates (digital) using asset_id.

Public APIs (read)

- GET /v1/assets/{asset_id} → metadata, transcript, thumbnails, provenance.
- GET /v1/races/{race_id}/overlaps → creative-hash, vendor-lattice, timing metrics.
- GET /v1/ie-filings?race_id=&since= → live IE ticker.
- GET /v1/vendors/{vendor_id}/family → affiliate graph.
- Bulk: daily ndjson dumps; monthly snapshots.

Filer APIs (write)

- POST /v1/ie-filings (commit, amount, schedule, vendor).
- POST /v1/assets (APB upload + media).
- POST /v1/corrections (amendment; original filing_id + delta).

Interchange & openness

- Open-source validators/CLIs; conformance suite with golden files.
- SPDX-like licensing for public dumps; deterministic IDs; clock skew tolerance ±5 min.

Privacy & safety

- Targeting disclosed at category/geo bins; no user-level IDs.
- Transcript profanity/PII masks for mail PDFs; original stored under regulator access.

14.2 Pilot jurisdictions and evaluation metrics

Phase 0 – Design lab (90 days)

- Partners: 1 state elections office + 2 broadcasters + 2 platforms + 3 watchdog orgs.
- Deliverables: schema freeze (APB v1), validator CLI, sandbox APIs, mock dashboards with synthetic data.

Phase 1 – Soft pilot (1 off-cycle city + 1 special election)

- Scope: voluntary APB uploads; state-hosted public library; live 48-hour IE ticker.
- Measures: latency (impression→public), validator pass rate, user comprehension (task-based tests).
- Success gates: median latency < 24h; >95% valid uploads; ≥3 newsroom embeds.

Phase 2 – State pilot (primary + general)

- Mandates: APB for all ads reaching ≥1,000 residents/24h; T+24 IE filings; vendor family registry.
- Enforcement: per-day fines for late files; safe harbor for corrected uploads within 48h.
- Added tools: whistleblower portal (bounties on discrete artifacts), rebuttal portal, red/yellow flags.

Phase 3 – Multi-state cohort (3–5 states)

- Interop: shared vendor family graph; federated API search; cross-state alerting.
- Research arm: pre-registered studies (event, RD, synthetic control) on reform effects.

Evaluation metrics (quant + qual)

- **Speed**: median time from first air/serve → public API; percent on-time (>90%).
- **Coverage**: share of ads with valid APB; % spend represented; number of unique entities compliant.

- **Detection**: count of yellow/red flags *pre-election*; corroboration rate across layers.
- **Behavior change**: cycle-over-cycle reduction in vendor overlaps and clone creatives; shift of spend to bright money.
- **Fairness**: reversal rate; median time to resolve rebuttals; stakeholder satisfaction (surveys).
- **Civic uptake**: dashboard DAUs, newsroom embeds, QR-card scans, average "Why this score?" dwell time.
- **Enforcement outcomes**: injunctions, disgorgements, voluntary ad pulls/vendor swaps.
- **Equity/access**: small-candidate upload success; helpdesk SLA; language/accessibility compliance.

Reporting cadence

- Weekly pilot briefs; election-night transparency report; 60-day postmortem with code/data annex.

14.3 Maintenance: audits, red-team drills, iteration

Operational auditing

- **Data audits (quarterly)**: random sample of assets; re-derive hashes; reconcile with broadcaster invoices and platform ad libraries.
- **Process audits (semiannual)**: review threshold changes, analyst actions, and flag histories; publish diffs and reasons.
- **Security audits (annual)**: third-party pen test of upload and public APIs; rotated signing keys; HSM checks.

Red-team drills (each cycle)

- **Blue team**: dashboard + enforcement staff. **Red team**: outside consultants and campaign veterans.
- Scenarios: paper-firewall breach via sub-vendors; creative clone with stock footage camouflage; timing-coupled bursts; conduit carousel using pop-up LLCs.
- Success criteria: time-to-detect (TDD), false-positive posture, quality of public narrative, and rebuttal handling.

Model & rule iteration

- **Threshold retrain** after each cycle using adjudicated truth set; publish model cards (inputs, limitations, drift).
- **Whitelist/blacklist refresh** for stock B-roll, party boilerplate, and common motifs.
- **Weight tuning** for Presumptive Coordination Score (PCS) with k-fold validation and out-of-sample checks.

Change management

- All schema/rule changes go through **Request for Comment (RFC)** with 14-day public window; version every artifact (APB v1→v1.1).
- Maintain **compatibility shims** for prior versions; deprecations with 2-cycle notice.

Helpdesk & education

- 24/7 election-week helpdesk; SLA: first response < 2h.
- Vendor/campaign training: quarterly webinars; "how to APB" starter kits; test sandboxes.
- Accessibility: multilingual guides, mobile-friendly uploaders, examples for micro-vendors.

Governance & ethics

- Independent oversight board (journalism, civil liberties, election law, ad-tech); conflict disclosures; annual public meeting.
- **Due-process ombuds** monitors rebuttal fairness; publishes reversal narratives and fixes.

Sustainability & funding

- Cost-share: state appropriations + platform carriage fees + philanthropic support; transparent budget.
- Open-source core; vendor ecosystem certified against conformance tests (badge program).

Sunset & review

- Every 2 cycles, run a "sunset review": keep, amend, or retire components based on impact metrics and public comment.
- Archive all data, code, and reports under an open license for research continuity.

North Star
Keep the **beam** regular (standards), the **shoreline** instrumented (APIs), and the **harbor** fair (due process). When upkeep is ritual—audits, drills, and humble iteration—the cheapest path stays the honest one.

15. Limits and Failure Modes

15.1 Adaptive adversaries and creative lawyering

Failure mode. Actors optimize against the rules you publish: swap to "issue ads," outsource to deniable affiliates, pre-bake creative libraries that evade hash matches, or route asks through general counsel memos that launder intent into "policy education."

How it shows up

- **Adversarial media edits:** micro-time-stretch VO, alter frame cadence, crop B-roll to defeat naive hashes.
- **Proxy orchestration:** "friends of" 501(c)s brief "the world," then only one IE moves within hours.
- **Legal choreography:** opinions that redefine "coordination" to a vanishing point ("public-materials reliance," "arm's-length affiliates").

Mitigations

- **Defense in depth:** require two orthogonal layers for public flags (e.g., creative + vendor, timing + messaging).
- **Model upgrades:** move from exact hashes to fused, adversarial-robust similarity (audio chroma + keyframe pHash + OCR n-grams).
- **Pattern statutes:** presumptions that combine signals (shared family vendor + ≤72h creative lineage) with reversible safe harbors (submit build logs).
- **Pre-commit design:** publish thresholds and windows but keep fusion weights and feature sets rotating (documented post-cycle).
- **Counter-lawyering playbook:** model rule comments and test-case memos ready for AGs/SECs to rebut "public materials" and "affiliate distance" defenses.

Residual risk
You won't catch all edge cases. Aim to make evasion slower, dearer, and riskier—not impossible.

15.2 Jurisdictional patchwork and platform resistance

Failure mode. Reforms land unevenly. One state runs T+24 with provenance; neighbors don't. Platforms/broadcasters balk at carrying compliance cost or claim preemption.

How it shows up

- **Spend flight:** money migrates to laxer states or to federal races.
- **Data gaps:** major platforms slow-roll APIs; small broadcasters miss filings; "walled gardens" cite privacy.
- **Preemption fog:** platforms claim federal supremacy against state transparency mandates.

Mitigations

- **Federated baseline:** common open schemas (APB v1), canonical vendor/family registries, and a **gateway API** that proxies to state libraries; one integration, many jurisdictions.
- **Carriage leverage:** condition state/broadcast carriage or ad credits on compliance; tiered fines per late-ad per day.
- **Procurement carrots:** fee discounts, priority inventory, and liability safe harbors for compliant platforms.
- **Interstate compacts:** states adopt mirror statutes and share enforcement; publish joint dashboards to blunt arbitrage.
- **Targeted scope:** disclose targeting at category/geo bins (not user-level) to neutralize privacy pretexts.
- **Litigation posture:** pre-draft AG opinions on non-preemption (transparency ≠ content regulation); amicus bench built in advance.

Residual risk
Patchwork persists. Prioritize high-population, high-spend states to cover most impressions; maintain a public "compliance map" so citizens and press see where the dark gaps remain.

15.3 Measuring impact without counterfactual omniscience

Failure mode. You cannot observe the world where reforms didn't exist. Critics call any change anecdote.

Measurement hazards

- **Attribution fog:** media cycles, scandals, macro shocks co-move with spend and votes.
- **Gaming the metrics:** actors adjust behavior to look "clean" on the dashboard while coordinating upstream.
- **Short windows:** elections are sparse events; statistical power is thin.

Mitigations

- **Multi-method triangulation:** pre-register analyses and combine Event Studies (5.1), RD in near-ties (5.2), and Synthetic Controls (5.3). Require convergence before strong claims.
- **Design metrics that are hard to game:** vendor-overlap coefficient, creative-lineage latency, disclosure latency—structural signals, not vibes.
- **Staggered adoption studies:** difference-in-differences where some states adopt reforms earlier; include placebo tests and event-time plots with pre-trend checks.
- **Instrumental variation:** exploit exogenous shocks (unexpected court orders, platform outages) as instruments for disclosure speed or ad visibility.
- **Public scorecard with error bars:** always publish CIs, bandwidth sensitivity, placebo ranks, and pre-registered analysis plans.
- **Behavioral proxies:** track shifts from dark to bright money, decline in ≤72h clone-creative deltas, and pre-election resolution rates of flags.

A compact impact frame
Let **EV_evasion = p·V − (C + K)** (Section 3.2). Successful reforms reduce **EV_evasion** by ↓p (RCV, vouchers), ↑C (provenance, vendor exclusivity), ↑K and ↓Δt (faster detection, penalties). Publish an annual estimate of each term's change with uncertainty and exemplars.

Residual humility
Declare what the system *can* evidence (structure, timing, lineage) and what it cannot (intent). Commit to post-cycle postmortems, open data/code, and sunset reviews that can shrink or retire components that don't move the needle.

16. Conclusion — Pricing the Workaround

16.1 Either the trick becomes impossible, or it grows teeth

This paper took a pragmatic view: if explicit quid-pro-quo is scarce and intent is deniable, then the venue for integrity is infrastructure. We engineered **cost** into the seams where covert coordination must leave traces—assets, vendors, timestamps, wires. With provenance bundles, exclusive-vendor rules, cooling-off windows, T+24 filings, adversarial forensics, and fast public dashboards, the evasions shrink to two outcomes:

1. **The trick becomes impossible.** Shared vendors can't straddle a race; staff can't carry tacit strategy across a 180-day firewall; creatives can't air without a signed pedigree; platforms can't hide buys off-ledger.
2. **The trick grows teeth.** If actors try anyway, their workaround becomes **self-incriminating**: hash lineages, timing couplers, vendor lattices, and funds-flow graphs light up *before* consequential votes; bounties and safe portals convert private risk into public receipts; penalties and reputational costs arrive on a short fuse.

We did not criminalize persuasion; we priced deception. We did not sermonize the sea; we **instrumented the shoreline**. The cheapest, safest strategy—staying truly independent and competing on persuasion—becomes the rational choice.

16.2 A citizen's, journalist's, and candidate's to-do list

For citizens

- **Check the race card.** Before you vote, open the dashboard for your district: look at the **Ad-Hash Overlap**, **Vendor Overlap**, and **48-Hour IE Ticker**.
- **Reward disclosure.** Prefer candidates and groups with active **Audit-Ready** badges; penalize those that refuse.
- **Use your vouchers (where available).** Deploy democracy vouchers or small donations *reactively* when a negative IE surge hits.
- **Report what you see.** If you work in production or ad ops, use the protected portal; if you're a viewer, submit screenshots and flight times—receipts matter.

For journalists

- **Embed the widgets.** Drop the **race dashboard** and **Pay-to-Vote Index** into stories so readers see evidence, not adjectives.

- **Ask for logs, not quotes.** Demand **APB** records, vendor lists (with **family IDs**), seat IDs, and change histories; publish rebuttals alongside receipts.
- **Pre-register scrutiny.** Announce, in advance, which high-spend races you'll monitor and which thresholds will trigger coverage.
- **Follow the patterns.** Cover **clone creatives**, **common agents**, **timed barrages**, and **conduit carousels** as *named patterns* with consistent rubrics.

For candidates and IEs (who want to win clean)

- **Sign the pledge.** Adopt the **Audit-Ready Attestation**: immutable creative logs, public vendor/sub lists, seat IDs, cooperation within 72 hours.
- **Choose exclusivity.** Hire vendors who commit to **race-exclusive** engagements across their entire corporate family and subs.
- **Build with provenance.** Stamp every asset with the **Ad-Provenance Bundle** at creation; insist that broadcasters/platforms verify and publish in T+24.
- **Design for rebuttal.** Maintain export trees, targeting schemas, and change logs so you can clear false flags in hours, not weeks.
- **Protect your juniors.** Bake the whistleblower safe clause into every NDA; instruct managers that retaliation is terminal to the campaign.

For policymakers and platforms

- **Codify the clocks.** Pass T+24 IE and ad-library mandates with open APIs; fund validators and helpdesks.
- **Close the seams.** Enact **exclusive-vendor/affiliate rules** and **cooling-off windows** with real penalties.
- **Standardize IDs.** Maintain the **Vendor Family Registry** and inter-state compacts so one integration works everywhere.
- **Publish the ledger.** Platforms and broadcasters: verify **APB**, expose **asset_id/seat_id**, and retain creatives for four years.

Closing note. A lighthouse is just a schedule, a lens, and a ledger turned toward the water. If we keep the **beam** regular (standards), the **shoreline** instrumented (APIs), and the **harbor** fair (due process), then covert coordination is either **impossible** or **expensive**—and in a democracy, that's the right kind of choice architecture.

Epilogue — The Lighthouse as a Mechanism

We did not chase the prettiest date. We built a lighthouse.

Nash taught that many fights are not won by better speeches but by better games. When everyone's payoffs change, everyone's *best reply* changes with them. Our program never tried to read souls or prove intent; it priced the workaround. Exclusive vendors, cooling-off windows, provenance tags, T+24 APIs, whistleblower bounties, rebuttal rites—each is a small gear. Together they rewire incentives so that "coordinate in secret" ceases to be the risk-dominant move.

Mechanism design, not moral suasion, is the throughline. We enumerated observables the system can actually see—assets, vendors, timestamps, wires—and bound them to clocks and public ledgers. That bundle becomes a correlation device: a shared signal that tells every player, *everyone else can see this in 24 hours too*. In such a game, the equilibrium tips. The cheapest path is the honest path; the clever path grows teeth.

Notice what we did *not* do. We didn't outlaw persuasion or confide in vibes. We didn't make truth depend on a confession that will never come. We replaced sermons with receipts, hunches with hashes, suspicion with schedules. The beam sweeps; the harbor watches; the payoffs settle.

If the analogy feels strange, it's only because civic life often mistakes eloquence for architecture. Lighthouses are architecture. They do not argue with fog; they measure it. They do not punish currents; they publish the tide. And because their rhythm is public and regular, ships adopt strategies that keep everyone alive. That is a Nash story told in salt and light: equilibria emerging from fixtures, not pleadings.

This is where we leave it. Not with a flourish, but with maintenance: audits, red-team drills, model cards, open APIs, and due process that dignifies rebuttal. Keep the beam regular. Keep the shoreline instrumented. Keep the harbor fair. If we do, two endings remain—either the trick becomes impossible, or it grows teeth and bites its owner.

In a politics that prizes the instant win, this can feel like a sidestep. It isn't. It is the shortest route to a stable peace: a design where honesty pays, evasion costs, and the prettiest date walks herself home—because the dance was set up that way from the start.

Appendix A: "Coordination Cost Raiser" — Field Checklist (One-Pager)

Goal: Make covert coordination slow, risky, and expensive—*before* consequential votes.

Who runs this: Campaign integrity lead + data analyst + counsel (+ media monitor on call)

Daily cadence (T+24 loop)

- Pull latest **IE filings** and **Ad Library** entries.
- Ingest creatives; compute **audio/video/text hashes**; OCR lower-thirds.
- Update **vendor lattice** (vendors, subs, corporate family IDs).
- Refresh **timing coupler** vs. committee markups/floor calendars.

Tripwires (trigger a case file)

- Creative: $S_{fused} \geq \tau$ (≥2 of audio/video/text) within **72h** of a campaign asset.
- Vendor: path length $d \leq 2$ between campaign ↔ IE via vendor family/sub.
- Timing: coupling ratio $R \geq 2$ around markup/floor windows.
- Messaging: semantic sim ≥ 0.85 inside 72h window for rare phrases.

Grab-list (evidence to capture immediately)

- Creative files, transcripts, checksums; export metadata (codec, LUT, timecodes).
- Buy sheets / impression aggregates by hour & DMA/ZIP.
- Vendor contracts, sub-lists, billing addresses, corporate family docs.
- Platform seat IDs, pixel IDs, audience taxonomies (category-level).
- Legislative timestamps; whip-sheet notes where lawful.

Escalation tiers

- **Yellow** (2 signals moderate): request attestations + build logs within **72h**.
- **Red** (3+ strong): publish dashboard card with slices; refer to regulator; preserve legal hold.

Do / Don't

- **Do** pre-register thresholds and windows; keep immutable logs.
- **Don't** accuse intent; narrate *structure & timing*.

- **Do** use whistleblower portal; protect juniors & vendors.
- **Don’t** overlook stock/party boilerplate—whitelist commons.

Contacts

- Regulator intake URL • Platform compliance emails • Broadcaster political-file desk • Pro bono counsel list

---

Appendix B: Data Schemas — Ad Hashes, Vendor Graphs, Disclosure APIs

B.1 Identifiers

- entity_id (campaign/IE/vendor/platform): ENT-<namespace>-<base32-8>
- family_id (corporate family): FAM-<base32-5>
- asset_id (creative content-address): adz:<sha256-16>
- seat_id (ad-tech buying seat): platform-issued string
- bundle_id (provenance bundle digest): apb:<sha256-16>
- filing_id (IE record, versioned): <uuid>:vN

B.2 Creative Hash Record (normalized)


```
{
 "asset_id": "adz:4C9F8A2B7D0E9C10",
 "ingested_ts": "2026-10-05T14:32:09Z",
 "media": {"kind":"tv","duration_s":30,"fps":29.97,"size":"1920x1080"},
 "hashes": {
  "audio_chroma": "bf1e…",
  "keyframe_phash": ["a1b2…","c3d4…"],
  "ocr_minhash": ["mha1…","mha2…"],
  "ssim_signature": "0.912"
 },
 "ocr": {"lower_thirds":["SEN. DOE — FAILED US ON HB123"],"disclaimer":"Paid for by …"},
```

```
  "text": {"asr_transcript":"…", "language":"en"},
  "provenance": {"bundle_id":"apb:9F2D…","uploader_entity_id":"ENT-IE-7X2K9MFT"},
  "links": {"filings":["FIL-…"], "platform_ids":["yt:…","fb:…"]}
}
```


B.3 Vendor Graph (bipartite lattice)


```
{
  "nodes": [
    {"id":"ENT-CAM-1ABCD123","type":"entity","role":"campaign","family_id":null},
    {"id":"ENT-IE-7X2K9MFT","type":"entity","role":"ie","family_id":null},
    {"id":"ENT-VND-5P8Q2ZJL","type":"vendor","role":"prime","family_id":"FAM-1R3T8"},
    {"id":"ENT-VND-9LMN4RTS","type":"vendor","role":"sub","family_id":"FAM-1R3T8"}
  ],
  "edges": [
    {"src":"ENT-CAM-1ABCD123","dst":"ENT-VND-5P8Q2ZJL","kind":"contract","since":"2026-08-01"},
    {"src":"ENT-IE-7X2K9MFT","dst":"ENT-VND-9LMN4RTS","kind":"subcontract","since":"2026-09-15"},
    {"src":"ENT-VND-5P8Q2ZJL","dst":"ENT-VND-9LMN4RTS","kind":"affiliate","evidence":"shared_officers"}
  ],
  "metrics": {"shortest_path_d": 2, "overlap_Ostar": 0.67}
}
```


B.4 IE Filing (write API, T+24)


```
{
  "filing_id": "FIL-3a2f…:v1",
  "filer": {"entity_id":"ENT-IE-7X2K9MFT","treasurer":"A. Smith"},
  "commit": {"amount_usd": 250000, "first_flight_ts": "2026-10-11T12:00:00Z"},
```

"race_id": "RACE-US-H-06-2026",

"vendor_id": "ENT-VND-5P8Q2ZJL",

"targets": {"geo":["DMA:803"],"categories":["age_25_54"]},

"assets": ["adz:4C9F…","adz:88AA…"],

"funding_since_last": [{"source":"donor_agg","amount_usd":250000}],

"signature": {"alg":"ed25519","key_id":"KEY-ENT-IE-7X2K9MFT","sig_b64":"…"}

}

B.5 Public Ad Library (read API)

- GET /v1/assets/{asset_id} → creative metadata + transcript + thumbnails + APB
- GET /v1/filings?race_id=&since= → streaming NDJSON of new/updated filings
- GET /v1/vendors/{vendor_id}/family → affiliate graph
- GET /v1/races/{race_id}/overlaps → {ad_hash_meter, vendor_overlap, timing_coupler}

Notes

- All endpoints support If-None-Match with ETags; time is UTC ISO-8601; dumps shipped as daily *.ndjson.zst.

Appendix C: Sample Statutes / Rules

C.1 Exclusive Vendor / Affiliate Rule

Section 1. Race-Exclusive Political Vendor Services.
(a) No vendor, subsidiary, affiliate, or subcontractor under common control shall provide strategic, creative, data, media buying, or production services to both (i) a candidate committee and (ii) an independent-expenditure entity influencing the same race in the same election cycle.
(b) “Common control” includes ≥20% ownership, shared officers, or shared back-office systems.
(c) Violations are subject to disgorgement, civil penalties up to 3× contract value, and disqualification for the remainder of the cycle.

C.2 Cooling-Off for Staff Transitions

Section 2. Cooling-Off Period.
Any individual with access to nonpublic strategic materials of a covered committee is prohibited

from providing services to an independent-expenditure entity in that race for **180 days**, and vice-versa. Willful violations render related expenditures coordinated in-kind contributions.

C.3 Provenance Tags & Public Ad Library (T+24)

Section 3. Ad Provenance & Disclosure.
(a) A covered communication may not be disseminated unless its media file embeds a signed **Ad-Provenance Bundle (APB)** containing buyer, vendor, family, asset checksum, first air/serve time, and targeting categories.
(b) Platforms and broadcasters shall validate APB at upload and publish an API-addressable record within **24 hours**.
(c) Failure to publish constitutes an unfair and deceptive act; each day per ad is a separate violation.

C.4 Rapid IE Reporting

Section 4. Independent Expenditure Rapid Reports.
Commitments ≥ $10,000 must be filed within **24 hours** of execution, including vendor and expected flights. Amend when commitments change by ≥10%.

C.5 Safe Harbors & Due Process

Section 5. Safe Harbor.
Good-faith filers correcting within **48 hours** after notice are exempt from penalties for the underlying late filing.
Section 6. Rebuttal Rights.
Accused entities shall receive a 72-hour rebuttal window via a secure portal to submit build logs, manifests, and attestations.

C.6 Private Right of Action

Section 7. Private Enforcement.
Any resident voter may seek injunctive relief and costs for violations of Sections 3–4.

Appendix D: Methods Notes (Templates)

D.1 Event Study (Spending Shock)

Spec

$Y_{it} = \alpha_i + \delta_t + \Sigma_{k=-K}^{+L} \beta_k * Shock_{i,t+k} + X_{it}'\gamma + \varepsilon_{it}$

- Units i = districts; t = days/weeks; K=L=3 typical.

- Pre-register **Shock** (first unanticipated IE flight).
- Plot {β_k} with 95% CIs; verify flat pre-trends.

Checklist

- Controls: national day FE, topic-news dummies.
- Placebos: pre-period fake shocks.
- Sensitivity: window ±7/±21 days; alternative outcomes (small-donor counts, search queries).

D.2 Local Linear RD (Near-Tie Races)

Spec

$Y = a + \tau * 1[r \geq 0] + b1 * r + b2 * r*1[r \geq 0] + \varepsilon \quad (|r| \leq h)$

- r = vote margin; choose bandwidth h via MSE-optimal + robustness.
- Tests: McCrary density; covariate balance.
- Report bias-corrected τ, multiple h, donut excluding recount outliers.

Augments

- Fuzzy RD where treatment (IE flood) jumps at cutoff imperfectly.
- RD-in-time for primary advancement.

D.3 Synthetic Control (Policy Flip)

Steps

1. Define treated member + flip date; compile predictors (past votes, ideology score, district covariates, donor mix).
2. Choose donor pool (no IE barrage); optimize weights to match pre-period.
3. Estimate post-gap $\Delta = Y_{treat} - Y_{synth}$.
4. Placebo reassignments to each donor; rank treated gap.

Diagnostics

- Pre-fit RMSPE threshold; leave-one-out stability.
- Augmented SC (ASC) if small-N.

Pre-Registration Template (abridged)

- Hypothesis, units, windows, outcomes
- Shock and threshold definitions
- Exclusion rules, covariates, FE structure
- Robustness plan, code repo URL, data release plan

Appendix E: “How to Keep a Lighthouse Honest” — Full Poem

How to Keep a Lighthouse Honest

The beam is a promise kept on a clock.
Not a sermon, not a shout—just light that arrives
when it said it would, and again, and again.
Keep the clock public. Let the sweep be seen.

Fog is not an argument you win.
It is weather. You measure it:
how far the bell carries,
how the buoy blinks and then disappears.
Say only what the fog allows; stamp it with the hour.

Currents do not confess.
They pull. They leave you somewhere else
than where your charts believed.
Mark their drift as numbers on a piling;
call the pull by name and publish the tide.

Wakes are signatures a hull cannot hide:
a frayed edge of foam, a shiver in the reeds,
the way borrowed B-roll repeats its grain.
Index every wake to the vessel that drew it,
and keep the index open to the town.

Build posts before you write poems.
Plant one at **timing**—T plus twenty-four, no later.
Plant one at **origin**—every lamp tagged with its maker.
Plant one at **contact**—no rope ties both boats this season.
Plant one at **redress**—a fast door, a fair ear.

String lines the harbor can see:
a line of sight (green, yellow, red with receipts),
a line of custody (hashes that outlast memory),
a line of care (a safe hand for the kid
who carried the reel between two shops).

Let ships argue policy in open water.
Your work is the shoreline—
instruments that price deceit,
ledgers that hum instead of preach.

If you do it right, two endings are left:
the trick becomes impossible,
or it grows teeth and bites its owner.
Either way, the beam stays steady.
Either way, the harbor learns to hum along—
a song of clocks, of measured fog,
of tides that tell the truth in inches,
and wakes that sign their authors in the dark.

References

Foundational U.S. cases & statutes

- **Buckley v. Valeo**, 424 U.S. 1 (1976). Summary & documents, Federal Election Commission. FEC.gov+1
- **Citizens United v. FEC**, 558 U.S. 310 (2010). Slip opinion & PDF. Legal Information Institute+1
- **SpeechNow.org v. FEC**, 599 F.3d 686 (D.C. Cir. 2010) (en banc). Opinion & FEC case page. FEC.gov+1
- **McDonnell v. United States**, 579 U.S. 550 (2016) (narrowing "official act"). Opinion & overview. Justia Law+1
- **Hobbs Act**, 18 U.S.C. § 1951 (extortion; "under color of official right"). Statute & DOJ manual. Legal Information Institute+1

U.S. regulations & administrative guidance

- **Coordinated communications** definition, 11 C.F.R. § 109.21 (payment, content, conduct prongs). Legal Information Institute
- FEC guidance: **Coordinated communications – candidate** (commercial vendor & conduct standards). FEC.gov
- FEC guidance: **Making independent expenditures** (what counts; limits and reports). FEC.gov
- FEC: **48-hour IE reports** (thresholds; repeated filings). FEC.gov
- FCC: **Political file / public inspection file** (broadcaster obligations; ad records). FCC Public Files+1

Transparency & platform/state policy references

- **Honest Ads Act** (proposed): current Congress bill pages & summaries. GovInfo+3Congress.gov+3Congress.gov+3
- Brennan Center explainer on the Honest Ads Act (policy rationale & scope). Brennan Center for Justice
- Washington State **commercial advertiser disclosure** statute & guidance (state ad-archive mandate). Washington State Legislature+1

- California/FPPC & advocacy materials on digital political ad archives (comparative notes). FPPC
- Academic/legal analyses of online political disclosure frameworks. NYU Law Review+1

Methods for causal inference used in Sections 5 & 15

- Abadie, A., Diamond, A., & Hainmueller, J. (2010). **Synthetic Control Methods for Comparative Case Studies**. *JASA*, 105(490), 493–505. (Working paper & publisher pages; later overviews.) NBER+2Taylor & Francis Online+2
- Imbens, G., & Lemieux, T. (2008). **Regression Discontinuity Designs: A Guide to Practice**. *Journal of Econometrics*. (NBER & publisher versions.) NBER+1
- Lee, D., & Lemieux, T. (2010). **Regression Discontinuity Designs in Economics**. *JEL*, 48(2), 281–355. Princeton University
- Angrist, J., & Pischke, J.-S. (2009). **Mostly Harmless Econometrics**. Princeton University Press. (Open PDFs/overviews). UCEMA+1
- Cunningham, S. (2021). **Causal Inference: The Mixtape**. Yale University Press (book site & open PDF). Causal Inference The Mixtape+1

Game theory & mechanism design touchstones for the Epilogue

- Nash, J. (1951). **Non-Cooperative Games**. *Annals of Mathematics*. (Open copies & university archive notes). Princeton University Library+2Computer Science Department+2
- Schelling, T. (1960). **The Strategy of Conflict**. Harvard University Press (library/open copies). Internet Archive+1
- Myerson, R. (1981). **Optimal Auction Design** (mechanism design classic). *Mathematics of Operations Research*. (JSTOR/open copies). Computer Science at Princeton+1
- Holt, C. (2004). **The Nash equilibrium: A perspective**. *PNAS*. PNAS

Background/context on U.S. campaign-finance evolution (non-authoritative but helpful)

- Overview journalism on the post-*Citizens United* landscape & state preemption fights. The Guardian+2The Guardian+2

Notes: Where multiple links are cited for a single item, the first is typically a primary-source PDF or official page; the second provides an accessible summary or alternate host.